Choreography Automata (TR)*

Franco Barbanera[1], Ivan Lanese[2], and Emilio Tuosto[3]
barba@dmi.unict.it, ivan.lanese@gmail.com, emilio.tuosto@gssi.it

[1] Dept. of Mathematics and Computer Science, University of Catania (Italy)
[2] Focus Team, University of Bologna/INRIA (Italy)
[3] Gran Sasso Science Institute (Italy) and University of Leicester (UK)

Abstract. Automata models are well-established in many areas of computer science and are supported by a wealth of theoretical results including a wide range of algorithms and techniques to specify and analyse systems. We introduce *choreography automata* for the choreographic modelling of communicating systems. The projection of a choreography automaton yields a system of *communicating finite-state machines*. We consider both the standard asynchronous semantics of communicating systems and a synchronous variant of it. For both, the projections of well-formed automata are proved to be live as well as lock- and deadlock-free.

1 Introduction

Choreographies are gaining momentum in the design and implementation of distributed applications also in the ICT industrial sector. This is witnessed by the effort of defining standards for specification languages such as WS-CDL [31] or BPMN [41] as well as the recognition of choreographies as suitable approaches to describe modern architectures such as microservices [12,2]. Choreographic approaches to the modelling, analysis, and programming of message-passing applications abound. For instance, in [35,5] abstract models have been applied to verify and debug BPMN specifications. Also, behavioural types have been proposed as suitable formalisations of choreographies [29] and for the analysis of properties such as liveness or deadlock freedom (e.g., [46,20] and the survey [30] to mention but few), while other approaches have considered syntax-free models [49]. At a programming level, choreographic programming has been explored in [36,40].

A distinguished trait of choreographies is the coexistence of two distinct but related views of a distributed system: the *global* and the *local* views. The former is an abstraction that yields a *holistic* description of the system. A global view indeed describes the coordination necessary among the various components of

* Research partly supported by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement No 778233, by the MIUR project PRIN 2017FTXR7S "IT-MaTTerS" (Methods and Tools for Trustworthy Smart Systems). and by the Piano Triennale Ricerca - UNICT 2016-19. The first and second authors have also been partially supported by INdAM as members of GNCS (Grup po Nazionale per il Calcolo Scientifico). The authors thanks the reviewers for their helpful comments. The first author thanks also M. Dezani for her everlasting support.

the system "altogether". In contrast, the local views specify the behaviour of the single components in "isolation".

In this paper we revisit the use of finite state automata to formally specify (and analyse) global views of message-passing systems, following an intuition similar to conversation protocols (CP) [26,16,27], a formalism where choreographies for asynchronous systems are described by means of Büchi automata. Our model, dubbed *choreography automata* (c-automata, for short), differs from CP in spite of some similarities. A comparison is given in § 6. The transitions of c-automata are labelled with *interactions*. As in most approaches, an interaction $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ states that *participant* A sends message m to participant B, which in turn receives it. For instance, consider the c-automaton

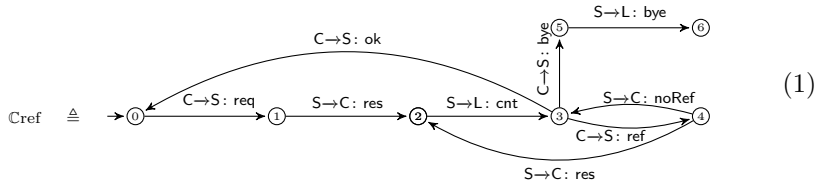

(1)

used to illustrate our model and as our working example through the paper. The c-automaton $\mathbb{C}\mathsf{ref}$ specifies the coordination among participants C, S, and L whereby a request req from client C is served by server S which replies with a message (of type) res and logs some meta-information cnt on a service L (e.g., for billing purposes). Client C may acknowledge a response of S (*i*) with an ok message to restart the protocol, or (*ii*) by requiring a refinement of the response with a ref message, or else (*iii*) by ending the protocol with a bye message which S forwards to L. In the second case, S sends C either a noRef message if no refinement is possible or another res (with the corresponding cnt to L).

Note that $\mathbb{C}\mathsf{ref}$ has nested as well as entangled loops. The support for entangled loops is a distinguished and expressive feature of c-automata, not present in many existing models of choreographies or multiparty session types (MST), and that we shall discuss in § 6.

We argue that c-automata provide a number of benefits. An advantage of c-automata is that finite state automata are well-known structures used both in theoretical and applied computer science. For instance, the c-automaton $\mathbb{C}\mathsf{ref}$ above can be easily understood by practitioners while retaining rigour. Another advantage is that c-automata are *syntax-independent*; they do not rely on complex linguistic constructs (such as the process algebraic constructs usually adopted in behavioural types). More crucially, we can re-use well-known results of the theory of automata and formal languages (e.g., we use determinisation and trace equivalence) as well as related algorithms. We discuss these advantages more extensively in § 6.

Choreographies enable a so called top-down approach whereby local views can be *projected* from the global view. Projections are expected to reflect the global specification without spoiling communication soundness (e.g., deadlock freedom, liveness, etc.). These results do not hold in general. In fact, global

views abstract away from "low level" peculiarities and projections may exhibit unintended behaviour.

The *realisability* of a global specification is a natural question to ask:

> Can global views such as $\mathbb{C}\mathsf{ref}$ be realised by distributed components C, S, and L coordinating with each other without intermediaries?

The answer to such question (obviously) depends on the communication infrastructure the distributed components use for the coordination. In fact, global views in general *abstractly* specify the coordination disregarding several details. For instance, the c-automaton $\mathbb{C}\mathsf{ref}$ in (1) is oblivious of the communication infrastructure used by the participants to coordinate with each other. Are the communications among C, S, and L synchronous or asynchronous? In the latter case, are messages received in their sending order? How is the sequencing reflected at the local level? For instance, should the messages that C sends from state 3 in (1) be sent after L receives the message cnt from S?

Tackling the realisability of global views is not a trivial endeavour. For instance, the recent analysis done in [46] highlights glitches in several projection operations of behavioural types. Also, some decidability results on the realisability of CPs [9], the only other automata-based choreographic setting that we are aware of, have been recently proved erroneous [24].

One would also like to understand whether the distributed components realising a choreography enjoy nice communication properties, e.g., will a component ready to engage in a communication eventually progress? Will a message sent by a participant eventually be received? We will consider such problems, showing that a set of conditions we define on c-automata do guarantee the choreography both to be realisable and to enjoy a number of relevant communication properties such as liveness and deadlock freedom.

Contributions & Structure. After a preliminary section (§ 2) recalling the main notions we deal with in the paper, in § 3 we formalise c-automata and their projections.We adopt *communicating systems* [13] (reviewed in § 2) for the local views of choreographies.

We consider both the case of synchronous and asynchronous communications for the local views. The projection from c-automata to communicating systems is defined in § 3 while in § 4 we define the class of *well-formed* c-automata for the synchronous case. There we show that, on well-formed c-automata, our notion of projection is correct (cf. Thm. 4.14) and guarantees liveness, lock- and deadlock-freedom in the synchronous semantics (cf. Thm. 4.15). In § 5 we generalise the above results to the case of asynchronous communications (cf. Thms. 5.6 and 5.7). Concluding remarks, related and future work are discussed in § 6. Additional material and complete proofs can be found in the appendixes.

Some interesting technical points are worth noticing. Firstly, most of our constructions and results rely on basic notions of formal languages and automata theory. This greatly simplifies the presentation and the proofs. The generalisation from synchronous to asynchronous communications requires only a mild

strengthening of our notion of well-formedness and no changes to c-automata or their projection. These are further advantages of the use of finite-state automata.

2 Preliminaries

A *Labelled Transition System* (LTS) is a tuple $A = \langle \mathbb{S}, s_0, \mathcal{L}, \rightarrow \rangle$ where

- $\mathbb{S}$ is a set of states (ranged over by $s, q, \ldots$) and $s_0 \in \mathbb{S}$ is the *initial state*;
- $\mathcal{L}$ is a finite set of labels (ranged over by $l, \lambda, \ldots$);
- $\rightarrow \subseteq \mathbb{S} \times (\mathcal{L} \cup \{ \varepsilon \}) \times \mathbb{S}$ is a set of transitions where $\varepsilon \notin \mathcal{L}$ is a distinguished label.

We define a Finite-State Automaton (FSA) as an LTS where $\mathbb{S}$ is finite. We use the usual notation $s_1 \xrightarrow{\lambda} s_2$ for the transition $(s_1, \lambda, s_2) \in \rightarrow$, and $s_1 \rightarrow s_2$ when there exists λ such that $s_1 \xrightarrow{\lambda} s_2$, as well as $\rightarrow^*$ for the reflexive and transitive closure of $\rightarrow$. The set of *reachable states* of A is $\mathcal{R}(A) = \{ s \mid s_0 \rightarrow^* s \}$.

Remark 2.1. Our definition of FSA omits the set of *accepting* states since we consider only FSAs where each state is accepting (which is the normal case in LTSs). We discuss this point further at the end of the paper. ◇

We recall standard notions on LTSs.

Definition 2.2 (Traces and Trace equivalence). *A* run *of an LTS* $A = \langle \mathbb{S}, s_0, \mathcal{L}, \rightarrow \rangle$ *is a (possibly empty) finite or infinite sequence of consecutive transitions starting at* s_0*. The* trace *(or* word*)* w *of a run* $(s_{i-1} \xrightarrow{\lambda_{i-1}} s_i)_{1 \leq i \leq n}$ *of* A *is the concatenation of the labels of the run (assume* $n = \infty$ *if the run is infinite), namely* $w = \lambda_0 \cdot \lambda_1 \cdots \lambda_n$*; label* ε*, as usual, denotes the identity element of concatenation; if the run is empty then* $w = \varepsilon$*.*

The language $\mathfrak{L}(A)$ *of* A *is the set of the traces of the runs of* A*. Two LTSs* A *and* B *are trace equivalent iff* $\mathfrak{L}(A) = \mathfrak{L}(B)$*. Also,* A accepts w *if* $w \in \mathfrak{L}(A)$*,* A accepts w from s *if* $w \in \mathfrak{L}(\langle \mathbb{S}, s, \mathcal{L}, \rightarrow \rangle)$*, and an* s*-run (resp.* s*-trace) of* A *is a run (resp. trace) of* $\langle \mathbb{S}, s, \mathcal{L}, \rightarrow \rangle$*.*

The notion of language in the definition above includes infinite words; this extends the standard notion of language accepted by an FSA. In particular, we consider an infinite word to be accepted by an FSA if each of its prefixes is accepted in the standard way. This is equivalent to look at an FSA both as a standard FSA and as a Büchi automaton where all the states are final.

Definition 2.3 (Deterministic LTSs). *An LTS* $A = \langle \mathbb{S}, s_0, \mathcal{L}, \rightarrow \rangle$ *is* deterministic *if*
– it is ε*-free, i.e. there is no transition of the form* $q \xrightarrow{\varepsilon} q'$*, and*
– whenever $q \xrightarrow{\lambda} q_1$ *and* $q \xrightarrow{\lambda} q_2$ *then* $q_1 = q_2$*.*
We denote the determinisation of A *(i.e. the translation of a nondeterministic LTS/FSA to a deterministic one) as* $\mathsf{det}(A)$ [4]*.*

[4] The result of $\mathsf{det}(A)$ may actually depend on the chosen algorithm, but that is irrelevant for our results.

We adopt communicating finite-state machines (CFSMs) [13] to model the local behaviour of systems of distributed components. The following definitions are borrowed from [13] and adapted to our context.

Let $\mathfrak{P}$ be a set of *participants* (or *roles*, ranged over by A, B, etc.) and $\mathcal{M}$ a set of *messages* (ranged over by m, n, etc.). We take $\mathfrak{P}$ and $\mathcal{M}$ disjoint.

Definition 2.4 (Communicating system). *A* communicating finite-state machine *(CFSM) is an FSA on the set*

$$\mathcal{L}_{act} = \{\mathsf{A}\,\mathsf{B}!\mathsf{m}, \mathsf{A}\,\mathsf{B}?\mathsf{m} \mid \mathsf{A}, \mathsf{B} \in \mathfrak{P}, \mathsf{m} \in \mathcal{M}\}$$

of actions. *The* subject *of an output (resp. input) action* $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ *(resp.* $\mathsf{A}\,\mathsf{B}?\mathsf{m}$*) is* A *(resp.* B*). A CFSM is* A-local *if all its transitions have subject* A*.*

A (communicating) system *is a map* $\mathsf{S} = (M_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *assigning an* A*-local CFSM* M_A *to each participant* $\mathsf{A} \in \mathcal{P}$ *such that* $\mathcal{P} \subseteq \mathfrak{P}$ *is finite and any participant occurring in a transition of* M_A *is in* $\mathcal{P}$*.*

Note that CFSMs may contain ε-transitions. However, projection (see Def. 3.3 below) yields ε-free CFSMs.

Besides being a well-known and widely adopted model, CFSMs are equipped with both synchronous and asynchronous semantics. This enables a uniform treatment of both communication models. The use of CFMSs is also helpful to compare c-automata with other models which are projected on CFSMs as well, such as global graphs [38] and some versions of global types [23].

The synchronous semantics of communicating systems is an LTS where labels are *interactions*:

$$\mathcal{L}_{\text{int}} \;=\; \{\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m} \mid \mathsf{A} \neq \mathsf{B} \in \mathfrak{P} \text{ and } \mathsf{m} \in \mathcal{M}\}$$

Definition 2.5 (Synchronous semantics). *Let* $\mathsf{S} = (M_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *be a communicating system where* $M_\mathsf{A} = \langle \mathbb{S}_\mathsf{A}, q_{0\mathsf{A}}, \mathcal{L}_{act}, \to_\mathsf{A}\rangle$ *for each participant* $\mathsf{A} \in \mathcal{P}$*. A* synchronous configuration of S *is a map* $s = (q_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *assigning a* local state $q_\mathsf{A} \in \mathbb{S}_\mathsf{A}$ *to each* $\mathsf{A} \in \mathcal{P}$*. We denote* q_A *by* $s(\mathsf{A})$ *and may denote* s *by* $\vec{q}$*.*

The synchronous semantics *of* S *is the transition system* $[\![\mathsf{S}]\!]^\mathsf{s} = \langle \mathbb{S}, \vec{q}_0, \mathcal{L}_{int}, \to\rangle$ *defined as follows*

- $\mathbb{S}$ *is the set of synchronous configurations of* S*, as defined above, and* $\vec{q}_0 = (q_{0\mathsf{A}})_{\mathsf{A}\in\mathcal{P}} \in \mathbb{S}$ *is the* initial *configuration*
- $\vec{q}_1 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} \vec{q}_2$ *if*
  1. $\vec{q}_1(\mathsf{A}) \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}}_\mathsf{A} \vec{q}_2(\mathsf{A})$ *and* $\vec{q}_1(\mathsf{B}) \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}}_\mathsf{B} \vec{q}_2(\mathsf{B})$*, and*
  2. *for all* $\mathsf{C} \neq \mathsf{A}, \mathsf{B}$*,* $\vec{q}_1(\mathsf{C}) = \vec{q}_2(\mathsf{C})$*.*

  In this case, we say that $\vec{q}_1(\mathsf{A}) \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}}_\mathsf{A} \vec{q}_2(\mathsf{A})$ *and* $\vec{q}_1(\mathsf{B}) \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}}_\mathsf{B} \vec{q}_2(\mathsf{B})$ *are* component transitions of $\vec{q}_1 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} \vec{q}_2$*.*
- $\vec{q}_1 \xrightarrow{\varepsilon} \vec{q}_2$ *if* $\vec{q}_1(\mathsf{A}) \xrightarrow{\varepsilon}_\mathsf{A} \vec{q}_2(\mathsf{A})$*, and for all* $\mathsf{B} \neq \mathsf{A}$*,* $\vec{q}_1(\mathsf{B}) = \vec{q}_2(\mathsf{B})$*.*

Note that ε-transitions in the semantics of a communicating system are induced by those of the constituent CFSMs. Also, $[\![\mathsf{S}]\!]^\mathsf{s}$ is finite; in fact, it is in general a non-deterministic automaton on the alphabet $\mathcal{L}_{\text{int}}$.

As one would expect, the notion of synchronous semantics is invariant under language equivalence of CFSMs.

Proposition 2.6. *Let* $\mathsf{S} = (M_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *and* $\mathsf{S}' = (M'_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *be two communicating systems. If* $\mathfrak{L}(M_\mathsf{A}) = \mathfrak{L}(M'_\mathsf{A})$ *for all* $\mathsf{A} \in \mathcal{P}$ *then* $\mathfrak{L}([\![\mathsf{S}]\!]^\mathsf{s}) = \mathfrak{L}([\![\mathsf{S}']\!]^\mathsf{s})$.

The asynchronous semantics of systems is defined in terms of transition systems which keep track of both the state of each machine and the content of unbounded FIFO queues b_{AB} which are associated to each channel $(\mathsf{A},\mathsf{B}) \in \mathcal{C}$, where $\mathcal{C} = \mathcal{P}\times\mathcal{P}\setminus\{(\mathsf{A},\mathsf{A}) \mid \mathsf{A}\in\mathcal{P}\}$. The queue b_{AB} is where M_A puts the messages to M_B and from which M_B consumes the messages from M_A. To avoid cumbersome parenthesis, we write $\mathsf{A}\,\mathsf{B} \in \mathcal{C}$ for $(\mathsf{A},\mathsf{B}) \in \mathcal{C}$.

Definition 2.7 (Asynchronous semantics). *Let* $\mathsf{S} = (M_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *be a communicating system where* $M_\mathsf{A} = \langle \mathbb{S}_\mathsf{A}, q_{0\mathsf{A}}, \mathcal{L}_{act}, \to_\mathsf{A}\rangle$ *for each participant* $\mathsf{A}\in\mathcal{P}$. *An* asynchronous configuration *of* S *is a pair* $s = \langle \vec{q}\ ;\ \vec{b}\rangle$ *where* $\vec{q} = (q_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *with* $q_\mathsf{A} \in Q_\mathsf{A}$ *and* $\vec{b} = (b_{\mathsf{AB}})_{\mathsf{AB}\in\mathcal{C}}$ *with* $b_{\mathsf{AB}} \in \mathcal{M}^*$; *we write* $s(\mathsf{A})$ *for* $\vec{q}(\mathsf{A})$ *and denote by* ε *the empty queue. The* asynchronous semantics *of* S *is the transition system* $[\![\mathsf{S}]\!]^\mathsf{a} = \langle \mathbb{S}, s_0, \mathcal{L}_{act}, \to\rangle$ *defined as follows*

- $\mathbb{S}$ *is the set of asynchronous configurations of* S *and* $s_0 = \langle \vec{q_0}\ ;\ \vec{b}\rangle$ *is the* initial *configuration where* $\vec{q_0} = (q_{0\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *and all the queues are empty.*
- $s \xrightarrow{l} s'$ *if* $s = \langle \vec{q}\ ;\ \vec{b}\rangle$, $s' = \langle \vec{q'}\ ;\ \vec{b'}\rangle$ *and either (1) or (2) below holds:*

1. $l = \mathsf{A}\,\mathsf{B}!\mathsf{m}$ *and* $s(\mathsf{A}) \xrightarrow{l}_\mathsf{A} s'(\mathsf{A})$ *and*

a. $s(\mathsf{C})' = s(\mathsf{C})$ $\forall \mathsf{C} \neq \mathsf{A} \in \mathcal{P}$ *and*
b. $b'_{\mathsf{A}\,\mathsf{B}} = b_{\mathsf{A}\,\mathsf{B}}.\mathsf{m}$ *and*
c. $b'_{\mathsf{C}\,\mathsf{D}} = b_{\mathsf{C}\,\mathsf{D}}$ *for all* $\mathsf{C}\,\mathsf{D} \in \mathcal{C}, \mathsf{C}\,\mathsf{D} \neq \mathsf{A}\,\mathsf{B}$

2. $l = \mathsf{A}\,\mathsf{B}?\mathsf{m}$ *and* $s(\mathsf{B}) \xrightarrow{l}_\mathsf{B} s'(\mathsf{B})$ *and*

a. $s'(\mathsf{C}) = s(\mathsf{C})$ $\forall \mathsf{C} \neq \mathsf{B} \in \mathcal{P}$ *and*
b. $b_{\mathsf{A}\,\mathsf{B}} = \mathsf{m}.b'_{\mathsf{A}\,\mathsf{B}}$ *and*
c. $b'_{\mathsf{C}\,\mathsf{D}} = b_{\mathsf{C}\,\mathsf{D}}$ *for all* $\mathsf{C}\,\mathsf{D} \in \mathcal{C}, \mathsf{C}\,\mathsf{D} \neq \mathsf{A}\,\mathsf{B}$

In the first (resp. second) case we say that $s(\mathsf{A}) \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}}_\mathsf{A} s'(\mathsf{A})$ *(resp.* $s(\mathsf{B}) \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}}_\mathsf{B} s'(\mathsf{B})$*) is a component transition of* $s \xrightarrow{\mathsf{A}\to\mathsf{B}:\,\mathsf{m}} s'$.

- $\langle \vec{q}\ ;\ \vec{b}\rangle \xrightarrow{\varepsilon} \langle \vec{q'}\ ;\ \vec{b'}\rangle$ *if* $\vec{q}(\mathsf{A}) \xrightarrow{\varepsilon} \vec{q'}(\mathsf{A})$ *for some* $\mathsf{A}\in\mathcal{P}$ *and for all* $\mathsf{B}\neq\mathsf{A}$, $\vec{q}(\mathsf{B}) = \vec{q'}(\mathsf{B})$, *and* $\vec{b} = \vec{b'}$.

State q_A keeps track of the state of the machine M_A and buffer b_{AB} keeps track of the messages sent from A to B (and not yet received by B). In a transition $s \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}} s'$, participant A adds message m in the queue of the channel $\mathsf{A}\,\mathsf{B}$ and symmetrically, in a transition $s \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}} s'$, participant B consumes message m from the top of the queue of the channel $\mathsf{A}\,\mathsf{B}$. In both cases, any machine or queue not involved in the transition is left unchanged.

The asynchronous semantics is also invariant under equivalence of CFSMs.

Proposition 2.8. *Let* $\mathsf{S} = (M_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *and* $\mathsf{S}' = (M'_\mathsf{A})_{\mathsf{A}\in\mathcal{P}}$ *be two communicating systems. If* $\mathfrak{L}(M_\mathsf{A}) = \mathfrak{L}(M'_\mathsf{A})$ *for all* $\mathsf{A} \in \mathcal{P}$ *then* $\mathfrak{L}([\![\mathsf{S}]\!]^\mathsf{a}) = \mathfrak{L}([\![\mathsf{S}']\!]^\mathsf{a})$.

For both the synchronous and the asynchronous semantics we restrict the attention to *fair runs*. An infinite run is fair if each transition which is continuously enabled is taken in a finite number of steps. A finite run is always fair.

We are interested in standard properties of communicating systems which we now recall. Definitions are alike in the synchronous and asynchronous semantics, hence, to avoid repetitions, below $[\![\_]\!]$ stands for $[\![\_]\!]^\mathsf{a}$ or $[\![\_]\!]^\mathsf{s}$.

Definition 2.9 (Communication properties). *Let* $\mathsf{S} = (M_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *be a communicating system.*

i) **Liveness***:* S *is* live *if for each configuration* $s \in \mathcal{R}(\llbracket \mathsf{S} \rrbracket)$*, each* $\mathsf{A} \in \mathcal{P}$ *and each outgoing transition* t *from* $s(\mathsf{A})$ *in* M_{A} *there exists a run of* $\llbracket \mathsf{S} \rrbracket$ *from* s *including a transition which has* t *as component transition.*

ii) **Lock freedom***: A configuration* $s \in \mathcal{R}(\llbracket \mathsf{S} \rrbracket)$ *is a* lock *if*
- *there is* $\mathsf{A} \in \mathcal{P}$ *with an outgoing transition* t *from* $s(\mathsf{A})$ *in* M_{A} *and*
- *there exists a run of* $\llbracket \mathsf{S} \rrbracket$ *starting from* s*, maximal w.r.t. prefix order, and containing no transition* t' *involving* A*.*

System S *is* lock-free *if for each* $s \in \mathcal{R}(\llbracket \mathsf{S} \rrbracket)$*,* s *is not a lock.*

iii) **Deadlock freedom***: A configuration* $s \in \mathcal{R}(\llbracket \mathsf{S} \rrbracket)$ *is a* deadlock *if*
- s *has no outgoing transitions in* $\llbracket \mathsf{S} \rrbracket$ *and*
- *there exists* $\mathsf{A} \in \mathcal{P}$ *such that* $s(\mathsf{A})$ *has an outgoing transition in* M_{A}*.*

System S *is* deadlock-free *if for each* $s \in \mathcal{R}(\llbracket \mathsf{S} \rrbracket)$*,* s *is not a deadlock.*

Liveness, as in [42], establishes the progress of communicating systems we are interested in. Lock freedom casts in our framework the idea that, similarly to [33,32], certain communications happen, whereas deadlock freedom extends the definition of deadlock in [19] to a setting which can be synchronous or asynchronous (as done also in [38,49]).

3 Choreography Automata

We introduce *choreography automata* (c-automata) as an expressive and flexible model of global specifications, following the styles of conversation protocols [27], choreographies [14,31,41], global graphs [49] and multiparty session types [17,28,30]. As customary in choreographic frameworks, we show how to project c-automata on local specifications. As anticipated, our projection yields a system of CFSMs formalising the local behaviour of the participants of a choreography.

C-automata (ranged over by $\mathbb{C}$A, $\mathbb{C}$B, etc.) are FSAs with labels in $\mathcal{L}_{\text{int}}$.

Definition 3.1 (Choreography automata). *A* choreography automaton *(c-automaton) is an FSA on the alphabet* $\mathcal{L}_{int}$*. Elements of* $\mathcal{L}^*_{int}$ *are* choreography words*, subsets of* $\mathcal{L}^*_{int}$ *are* choreography languages*.*

Remark 3.2. Def. 3.1 admits non-deterministic c-automata. This does not increase the expressiveness of our framework. In fact, (*i*) the notions that we use for our results rely on traces and (*ii*) our projection operation (cf. Def. 3.3) is insensitive to non-determinism (cf. Prop. 3.6). Non-deterministic specifications are however desirable since they are easier to attain for the designer. ◇

Given a c-automaton, our projection operation builds the corresponding communicating system consisting of the set of projections of the c-automaton on each participant, each projection yelding a CFSM. Hereafter, $\mathcal{P} \subseteq \mathfrak{P}$ is the set of participants of c-automata; note that $\mathcal{P}$ is necessarily finite.

Definition 3.3 (Automata Projection). *The* projection on A *of a transition* $t = q \xrightarrow{\lambda} q'$ *of a c-automaton, written* $t{\downarrow_{\mathsf{A}}}$*, is defined by:*

$$t{\downarrow_{\mathsf{A}}} = \begin{cases} q \xrightarrow{\mathsf{A}\,\mathsf{C}!\mathsf{m}} q' & \text{if } \lambda = \mathsf{B}{\to}\mathsf{C}\colon \mathsf{m} \quad \text{and} \quad \mathsf{B} = \mathsf{A} \\ q \xrightarrow{\mathsf{B}\,\mathsf{A}?\mathsf{m}} q' & \text{if } \lambda = \mathsf{B}{\to}\mathsf{C}\colon \mathsf{m} \quad \text{and} \quad \mathsf{C} = \mathsf{A} \\ q \xrightarrow{\varepsilon} q' & \text{otherwise} \end{cases}$$

The projection *of a c-automaton* $\mathbb{C}\mathrm{A} = \langle \mathbb{S}, q_0, \mathcal{L}_{int}, \to \rangle$ *on a participant* $\mathsf{A} \in \mathcal{P}$*, denoted* $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$*, is obtained by determinising and minimising up-to-language equivalence the* intermediate *automaton*

$$A_{\mathsf{A}} = \langle \mathbb{S}, q_0, \mathcal{L}_{act} \cup \{\,\varepsilon\,\}, \{\,(q \xrightarrow{\lambda} q'){\downarrow_{\mathsf{A}}} \mid q \xrightarrow{\lambda} q'\,\} \rangle$$

The projection of $\mathbb{C}\mathrm{A}$*, written* $\mathbb{C}\mathrm{A}\downarrow$*, is the communicating system* $(\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$*. The projection function trivially extends to choreography words and languages.*

The projection defined above, apart for determinisation and minimisation, is essentially homomorphic, as most of the projections in the literature. Other approaches such as [44,25] add hidden communications to be able to deal with larger classes of choreographies. We prefer the former approach for its simplicity. Hidden communications can however be added directly at the choreographic level as proposed in [37].

It is a simple observation that the projection on A of $\mathbb{C}\mathrm{A}$ is A-local, deterministic and hence ε-free. Thanks to the properties of determinisation and minimisation (as, e.g., in the partition refinement algorithm [43]), the states of $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$ are sets of states of $\mathbb{C}\mathrm{A}$.

Example 3.4 (Projections of $\mathbb{C}\mathrm{ref}$*).* The projections of our working example are

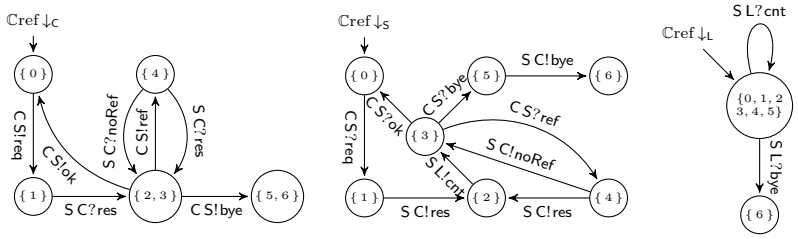

For instance, $\mathbb{C}\mathrm{ref}\downarrow_{\mathsf{C}}$ is obtained by determinising (minimisation is the identity in this case) the following intermediate automaton obtained as described in Def. 3.3.

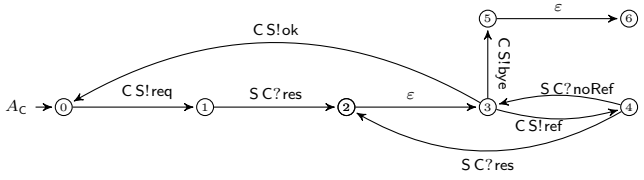

◇

The following proposition relates the language of the projection with the language of the original automaton.

Proposition 3.5. *For all c-automata* $\mathbb{C}\mathrm{A}$ *and* $\mathsf{A} \in \mathcal{P}$, $\mathfrak{L}(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{A}}}) = \mathfrak{L}(\mathbb{C}\mathrm{A}){\downarrow_{\mathsf{A}}}$.

The projection operation is well-behaved with respect to trace equivalence.

Proposition 3.6. *If* $\mathbb{C}\mathrm{A}$ *and* $\mathbb{C}\mathrm{A}'$ *are trace-equivalent c-automata then* $\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{A}}}$ *and* $\mathbb{C}\mathrm{A}'{\downarrow_{\mathsf{A}}}$ *are isomorphic for each participant* $\mathsf{A} \in \mathcal{P}$.

4 Well-formed Choreography Automata

To ensure that the communicating system obtained by projection of a c-automaton is well-behaved, some conditions are necessary. Since the conditions depend on the used communication infrastructure, we consider first synchronous communication, leaving to § 5 the case of asynchronous communication.

Definition 4.1 (Concurrent transitions). *Let* $\mathbb{C}\mathrm{A} = \langle \mathbb{S}, q_0, \mathcal{L}, \rightarrow \rangle$. *Two transitions* $q \xrightarrow{l_1} q_1$ *and* $q \xrightarrow{l_2} q_2$ *are concurrent iff there is a state* $q' \in \mathbb{S}$ *and transitions* $q_1 \xrightarrow{l_2} q'$ *and* $q_2 \xrightarrow{l_1} q'$.

Well-branchedness (cf. Def. 4.6) is a key notion which intuitively states that each participant is aware of the choices made in the choreography when its behaviour depends on those choices. The awareness of choice is checked on *spans*, namely pairs of runs that may constitute alternative branches of choices. Spans are formalised building on the notion of *candidate branch* which, in turn, is defined in terms of *pre-candidate branch*.

Definition 4.2 (Candidate q-branch). *Let* q *be a state of a c-automaton* $\mathbb{C}\mathrm{A}$. *A* pre-candidate q-branch *of* $\mathbb{C}\mathrm{A}$ *is a* q*-run of* $\mathbb{C}\mathrm{A}$ *such that each cycle has at most one occurrence within the whole run (i.e. any subsequence of the form* $q \rightarrow \ldots \rightarrow q$*, where* q *occurs only at the beginning and at the end of the subsequence, is not present more than once in the run). A* candidate q-branch *is a maximal pre-candidate* q*-branch with respect to prefix order.*

We often refer to a (pre-)candidate q-branch simply as "(pre-)candidate of q". Due to the condition about cycles in Def. 4.2, the following holds trivially.

Fact 1. *Given a state* q *of a c-automaton* $\mathbb{C}\mathrm{A}$*, the set of its pre-canditates is finite, and so is, a fortiori, that of its candidates.*

Example 4.3 ((Pre-)candidate branches in $\mathbb{C}$ref*).* The sequences in Fig. 1 are runs of the c-automaton of our working example. They are all pre-candidates of either 3 or 4, but run π_e, which is not a pre-candidate of 4 since the cycle $4 - 3 - 4$ occurs twice. Runs π_b and π_d are also candidates of 3, being maximal pre-candidates with respect to prefix order. ◇

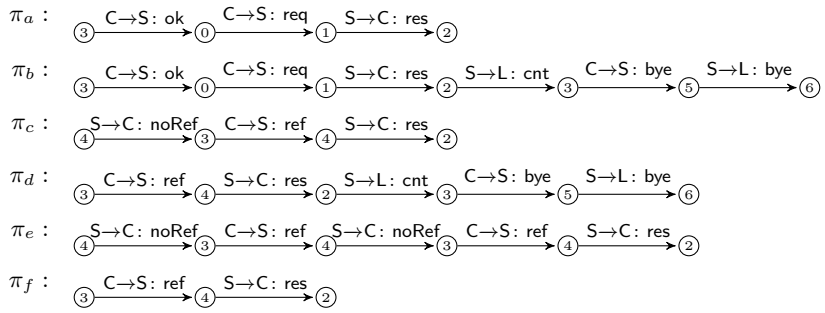
π_a :
3 C→S: ok 0 C→S: req 1 S→C: res 2
π_b :
3 C→S: ok 0 C→S: req 1 S→C: res 2 S→L: cnt 3 C→S: bye 5 S→L: bye 6
π_c :
4 S→C: noRef 3 C→S: ref 4 S→C: res 2
π_d :
3 C→S: ref 4 S→C: res 2 S→L: cnt 3 C→S: bye 5 S→L: bye 6
π_e :
4 S→C: noRef 3 C→S: ref 4 S→C: noRef 3 C→S: ref 4 S→C: res 2
π_f :
3 C→S: ref 4 S→C: res 2

Fig. 1. Runs of $\mathbb{C}$ref.

Definition 4.4 (q-span). *Given a state q of a c-automaton $\mathbb{C}\mathrm{A}$, a pair (σ, σ') of pre-candidate q-branches of $\mathbb{C}\mathrm{A}$ is a* q-span *if σ and σ' are*

- *either cofinal, with no common node but q and the last one;*
- *or candidate q-branches with no common node but q;*
- *or a candidate q-branch and a loop on q with no other common nodes.*

A participant $\mathsf{A} \in \mathcal{P}$ chooses *at a q-span (σ, σ') if the first transition of both σ and σ' has A as sender.*

Example 4.5 (Spans of $\mathbb{C}$ref*).* The states with spans of our working example are 3 and 4. A span from 3 is (π_a, π_f), where π_a and π_f are as in Fig. 1. Indeed, π_a and π_f are cofinal (in 2) pre-candidates of 3 with no common states but 3 and 2. Participant C chooses at (π_a, π_f). The pair (π_b, π_d), instead, is not a span from 3, since π_b and π_d are maximal, but share other nodes than 3. ◇

Intuitively, a choice is well-branched when the participants other than the one opting for alternative runs either behave uniformly in each branch, or can ascertain which branch has been chosen from the messages they receive.

Definition 4.6 (Well-branchedness). *A c-automaton $\mathbb{C}\mathrm{A}$ is* well-branched *if for each state q in $\mathsf{det}(\mathbb{C}\mathrm{A})$ and $\mathsf{A} \in \mathcal{P}$ sender in a transition from q, all of the following conditions must hold:*

(1) all transitions from q involving A, have sender A;
(2) for each transition t from q whose sender is not A and each transition t' from q whose sender is A, t and t' are concurrent
(3) for each q-span (σ, σ') where A chooses at and each participant $\mathsf{B} \neq \mathsf{A} \in \mathcal{P}$, the first pair of different labels on the runs $\sigma\!\downarrow_{\mathsf{B}}$ and $\sigma'\!\downarrow_{\mathsf{B}}$ (if any) is of the form $(\mathsf{C}\,\mathsf{B}?\mathsf{m}, \mathsf{D}\,\mathsf{B}?\mathsf{n})$ with $\mathsf{C} \neq \mathsf{D}$ or $\mathsf{m} \neq \mathsf{n}$.

We dub A a selector at q.

In the above definition loops are taken into account in item (3) since the notion of span is defined in terms of candidate branch. The latter is a maximal run where cycles can be considered at most once, as shown in Example 4.3.

In case of a nondeterministic c-automaton, the conditions of Def. 4.6 are checked after the c-automaton has been determinised. In fact, recalling Remark 3.2, we consider properties of languages of c-automata, and determinisation, as well as minimisation, of FSA preserve languages. Also, both operations preserve the system resulting from projection (cf. Prop. 3.6). (Observe that here we exploit classical results of automata theory.) Also, by Fact 1 and the obvious decidability of the conditions of Defs. 4.4 and 4.6 we get

Fact 2. *Well-branchedness is a decidable property.*

Example 4.7 (Well-branchedness of $\mathbb{C}$ref*).* All the states of $\mathbb{C}$ref satisfy the conditions of Def. 4.6; the only non-trivial cases are states 3 and 4. Condition (1) holds for C, which is the selector of the choice at 3, and for S, which is the selector of the choice at 4; condition (2) holds vacuously, and condition (3) holds for both S and L in all the spans from 3 and from 4. For instance, in the span (π_a, π_f) from 3, described in Example 4.5, the first actions of S on π_a and π_f are the inputs from C which have different messages, whereas, for what concerns L, the condition holds vacuously. As a matter of fact, since π_a and π_f are cofinal in 2, the well-branchedness conditions on state 2 do guarantee L to behave properly afterwards, independently on whether π_a or π_f have been followed before. ◇

Condition (2), vacuously true in our working example, is needed when multiple participants act as sender in the same state: this ensures that the only possibility is that actions of different participants are concurrent so that possible choices at a state are not affected by independent behaviour.

We add a further condition to rule out c-automata having consecutive transitions involving disjoint participants and not actually concurrent.

Definition 4.8 (Well-sequencedness). *A c-automaton* $\mathbb{C}$A *is* well-sequenced *if for each two consecutive transitions* $q \xrightarrow{\mathsf{A}\longrightarrow\mathsf{B}\colon \mathsf{m}} q' \xrightarrow{\mathsf{C}\longrightarrow\mathsf{D}\colon \mathsf{n}} q''$ *either*

- *they share a participant, that is* $\{\mathsf{A},\mathsf{B}\} \cap \{\mathsf{C},\mathsf{D}\} \neq \emptyset$*, or*
- *they are concurrent, i.e. there is* q''' *such that* $q \xrightarrow{\mathsf{C}\longrightarrow\mathsf{D}\colon \mathsf{n}} q''' \xrightarrow{\mathsf{A}\longrightarrow\mathsf{B}\colon \mathsf{m}} q''$*.*

Notice that, by finiteness of the transition relation of c-automata, we get

Fact 3. *Well-sequencedness is a decidable property.*

Notation. For the sake of readability, a well-sequenced c-automaton can be represented by omitting, for each diamond, two of its consecutive transitions. We call such representation *compact*. Notice that, given a compact representation, it is always possible to recover the original c-automaton. So far and hereafter we assume that all c-automata are compactly represented.

Example 4.9 (Well-sequencedness of $\mathbb{C}$ref*).* It is not difficult to check that $\mathbb{C}$ref is well-sequenced because the first condition of Def. 4.8 holds for any pair of consecutive transitions in $\mathbb{C}$ref. ◇

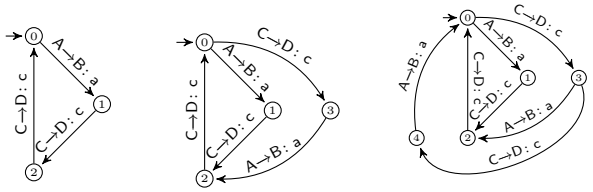

(a) (b) (c)

Fig. 2. Failure of well-sequencedness completion.

Well-sequencedness is necessary to establish a precise correspondence between the language of a c-automaton and of its projection (cf. Thm. 4.14 and the discussion following it).

Remark 4.10. We show that not all c-automata can be "completed" to well-sequenced ones. Consider the c-automaton of Fig. 2(a), which is not well-sequenced because of the transitions from state 0 to state 1 and from state 1 to 2. By "completing the diamond" for such transitions (i.e., by adding the new state 3 and the transitions $0 \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{c}} 3$ and $3 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{a}} 2$) we obtain the c-automaton of Fig. 2(b). This is still not well sequenced, because of the transitions $3 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{a}} 2$ and $2 \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{c}} 0$. So we try to make it well-sequenced by completing the diamond once again and obtain the c-automaton of Fig. 2(c). The resulting c-automaton is still not well-sequenced, because of the transitions $4 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{a}} 0$ and $0 \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{c}} 3$. Again a vain attempt, because of the transitions $5 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{a}} 3$ and $3 \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{c}} 4$. It is immediate to check that we could go on indefinitely.

It is impossible to complete the initial c-automaton since the intended completed automaton should generate a non-regular language (since it should generate strings with a number of $\mathsf{C}\to\mathsf{D}\colon \mathsf{c}$ interactions which is, roughly, double of the number of $\mathsf{A}\to\mathsf{B}\colon \mathsf{a}$ interactions). It would hence be interesting to know whether, in case the expected completed interaction language of a c-automaton is regular and prefix-closed, it is possible to generate it also by means of a well-sequenced c-automaton. It would be also interesting to establish a condition on cycles (if any) that guarantees the effectiveness of the completion of a c-automaton. We leave these questions for future work. ◇

We show a closure property of the languages of well-sequenced c-automata.

Definition 4.11 (Concurrency closure). *The* swap *relation on choreography words is the smallest equivalence relation* $\_ \sim \_$ *satisfying*

$$w(\mathsf{A}\to\mathsf{B}\colon \mathsf{m})(\mathsf{C}\to\mathsf{D}\colon \mathsf{n})w' \sim w(\mathsf{C}\to\mathsf{D}\colon \mathsf{n})(\mathsf{A}\to\mathsf{B}\colon \mathsf{m})w'$$

where $\{\mathsf{A},\mathsf{B}\} \cap \{\mathsf{C},\mathsf{D}\} = \emptyset$*. Given a choreography language* $\mathcal{L}$

$$close(\mathcal{L}) = \{\, w \in \mathcal{L}_{int} \;\big|\; \exists w' \in \mathcal{L}.\ w \sim w' \,\}$$

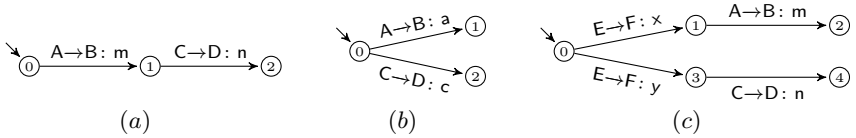

Fig. 3. Well-formedness is necessary for Thm. 4.14.

is the concurrency closure *of* $\mathcal{L}$.

The above relation is reminiscent of the *swapping* relation introduced in [18], with similar aims.

Proposition 4.12. *Let* $\mathbb{C}\mathrm{A}$ *be a well-sequenced c-automaton. Then* $\mathfrak{L}(\mathbb{C}\mathrm{A})$ *is concurrency closed, i.e.* $\mathfrak{L}(\mathbb{C}\mathrm{A}) = \mathit{close}(\mathfrak{L}(\mathbb{C}\mathrm{A}))$.

Notice that the converse of the above proposition does not hold in general. In fact, consider the following c-automaton

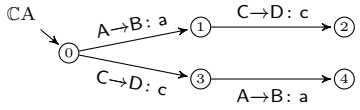

we can check that $\mathfrak{L}(\mathbb{C}\mathrm{A}) = \mathit{close}(\mathfrak{L}(\mathbb{C}\mathrm{A}))$ but $\mathbb{C}\mathrm{A}$ is not well-sequenced.

The notion of well-formedness below sums up the requirements needed in order for a c-automaton to be projected to a well-behaved communicating system.

Definition 4.13 (Well-formedness). *A c-automaton is* well-formed *if it is both well-branched and well-sequenced.*

The next result in Thm. 4.14 establishes that the language of a well-formed c-automaton coincides with the language of the communicating system obtained by projection. This provides a correctness criterion for our projection operation.

Theorem 4.14. $\mathfrak{L}(\mathbb{C}\mathrm{A}) = \mathfrak{L}(\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathsf{s}})$ *for any well-formed c-automaton* $\mathbb{C}\mathrm{A}$.

Notice that well-formedness is a necessary condition for the theorem above. It is in fact easy to check that

$$(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{m})(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{n}) \in \mathfrak{L}(\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathsf{s}}) \quad \text{and} \quad (\mathsf{C}{\to}\mathsf{D}\colon \mathsf{m})(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{n}) \notin \mathfrak{L}(\mathbb{C}\mathrm{A})$$

when $\mathbb{C}\mathrm{A}$ is one of the c-automata (a), (b) or (c) of Fig. 3. In particular, (a) is not well-sequenced whereas (b) and (c) are not well-branched: for (b), item (2) of well-branchedness (Def. 4.6) does not hold; (c) instead violates item (3).

We can now show that the projections of well-formed choreography automata enjoy the communication properties of Def. 2.9.

Theorem 4.15. *Given a well-formed c-automaton* $\mathbb{C}\mathrm{A}$*, its projection* $(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{A}}})_{\mathsf{A}\in\mathcal{P}}$ *is live, lock-free, and deadlock-free with respect to the synchronous semantics.*

5 Asynchronous Communications

We now transfer the results of the previous sections to the asynchronous semantics of communicating systems (Def. 2.7). Remarkably, the semantics does not affect the definition of c-automata (and of projections) since it is independent of the communication model. Hence, any result depending only on the definition of c-automata still holds. Well-sequencedness instead needs updating.

Definition 5.1 (Asynchronous well-sequencedness). *A c-automaton is* asynchronously well-sequenced *if for each two consecutive transitions* $q \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q' \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{n}} q''$ *either*

- *the sender of the second transition occurs in the first one, that is* $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$,
- *or they are concurrent, i.e. there is* q''' *such that* $q \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{n}} q''' \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q''$.

Asynchronous well-sequencedness (Def. 4.8) implies the synchronous one. Indeed, asynchronous well-sequencedness requires either two transitions to be concurrent or that the sender of the second transition occurs in the first one. The latter condition is weaker than having disjoint participants as required in the synchronous case.

Note that our working example is well-sequenced but not asynchronously well-sequenced (because e.g., of transitions $2 \xrightarrow{\mathsf{S}\to\mathsf{L}\colon \mathsf{cnt}} 3 \xrightarrow{\mathsf{C}\to\mathsf{S}\colon \mathsf{bye}} 5$). Thus, we now consider it as the *compact* representation of the actual c-automaton according to Notation on page 11.

Unlike well-sequencedness, the notion of well-branchedness has not to be changed in case asynchronous communications are considered. So, in the asyncronous setting, we define *asynchronous well-formedness* as the conjunction of asynchronous well-sequencedness (Def. 5.1) and well-branchedness (Def. 4.6).

The correspondence result between the semantics of a c-automaton and of its projection requires to decide which actions to observe on the projection. Indeed, in a c-automaton, each interaction is seen as an atomic event, while in the asynchronous semantics of communicating systems each interaction corresponds to two events: a sending event and a receving event. We opt to observe sending events only because (internal) choices are determined by sending events. This decision also plays well with the notion of well-branchedness, where most of the conditions concern sender participants. Other possible options are discussed in [36], in a process algebraic setting. This idea is formalised by *sender traces*.

Definition 5.2 (Sender traces). *The* sender traces $\mathfrak{L}^!(\mathsf{S})$ *of a communicating system* S *are obtained from its asynchronous traces by replacing each output label* $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ *with* $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ *and each input label* $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ *with* ε.

The modification of well-sequencedness for the asynchronous case does imply that we need to "update" the definition of concurrency closure as well.

Definition 5.3 (Asynchronous concurrency closure). *The* asynchronous swap *relation on choreography words is the smallest pre-order* $\_ \leq \_$ *satisfying*

$$w(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m})(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{n})w' \leq w(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{n})(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m})w' \qquad \text{where } \mathsf{C} \notin \{\mathsf{A},\mathsf{B}\}.$$

The downward closure of a choreography language $\mathcal{L}$ *with respect to* $\_ \leq \_$

$$close_{\mathrm{a}}(\mathcal{L}) = \{\, w \in \mathcal{L}_{int} \mid \exists w' \in \mathcal{L}.\ w \leq w' \,\}$$

is the asynchronous concurrency closure *of* $\mathcal{L}$.

The condition for asynchronous concurrency closure is weaker than the one in the synchronous case. This is due to the fact that sender-traces must be closed under asynchronous concurrency (cf. Lemma 5.4 below), so to guarantee that the traces of an automaton do coincide with the sender-traces of its projection (Thm. 5.6 below). We discuss such a necessity with an example after Thm. 5.6.

Lemma 5.4. *Let* $\mathbb{C}\mathrm{A}$ *be a c-automaton. Then* $\mathfrak{L}^{!}(\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathrm{a}}) = close_{\mathrm{a}}(\mathfrak{L}^{!}(\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathrm{a}}))$.

We now proceed to prove the correctness of projection for asynchronous systems. We will reduce it to the corresponding result for synchronous systems (Thm. 4.14). This is done by showing that all asynchronous runs are pairable (see below), that is they can be put in a suitable normal form which directly corresponds to a synchronous run. Notably, such a result is false for c-automata which are not asynchronously well-formed.

Definition 5.5 (Pairable runs). *Let* $\mathbb{C}\mathrm{A}$ *be a c-automaton. A run* σ *in* $\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathrm{a}}$ *is paired into a run* σ' *in* $\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathrm{a}}$ *iff they are coinitial, produce the same sender trace, and each output* $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ *in* σ' *is immediately followed by the corresponding input* $\mathsf{A}\,\mathsf{B}?\mathsf{m}$. *A run* σ *is* pairable *if it is paired into a run* σ'.

Theorem 5.6. *Let* $\mathbb{C}\mathrm{A}$ *be an asynchronously well-formed c-automaton.*

$$\mathfrak{L}^{!}(\llbracket \mathbb{C}\mathrm{A}{\downarrow} \rrbracket^{\mathrm{a}}) = \mathfrak{L}(\mathbb{C}\mathrm{A})$$

Similarly to Thm. 4.14, asynchronous well-formedness is a necessary condition for Thm. 5.6. Examples (*b*) and (*c*) of Fig. 3 work the same also for the asynchronous case, since we do not changed the definition of well-branchedness. We changed instead the definition of well-sequencedness to a stricter version and the c-automaton (*a*) of Fig. 3 is hence not enough to show the necessity of asyncronous well-sequencedness; this can however be easily done using the following c-automaton which is well-sequenced but not asynchronouly well-sequenced

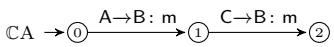

Since outputs of asynchronous CFSMs can always be fired, there is a run of the projected system beginning with $\mathsf{C}\,\mathsf{B}!\mathsf{m}$ and producing the sender trace $(\mathsf{C}{\to}\mathsf{B}\colon \mathsf{m})(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}) \in \mathfrak{L}^{!}(\llbracket \mathbb{C}\mathrm{A} \rrbracket^{\mathrm{a}})$ which trivially does not belong to $\mathfrak{L}(\mathbb{C}\mathrm{A})$ because the interactions cannot be swapped (cf. Def. 5.3).

The communication properties for projected systems can also be obtained.

Theorem 5.7. *Given an asynchronous well-formed c-automaton* $\mathbb{C}\mathrm{A}$*, its projection* $(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{A}}})_{\mathsf{A}\in\mathcal{P}}$ *is live, lock-free, and deadlock-free with respect to the asyncronous semantics.*

6 Conclusion, Related Work and Future Work

We introduced a model of choreographies based on FSAs whose transitions are labelled by interactions. We showed relevant results both for a synchronous and an asynchronous underlying communication infrastructure. We established a correspondence between the language of an automaton and the one of its projection, as well as proofs of liveness, lock, and deadlock freedom for the latter.

The adoption of an automata-based model brings in two main benefits. Firstly, the constructions that we provided are based on set-theoretic notions and are *syntax-independent*. This contrasts with syntax-driven models (such as behavioural type systems [30]) where expressiveness may be limited and definitions may be more complex due to syntactic reasons. E.g., the example in § 1 cannot be modelled in many behavioural type systems since entangled loops cannot be represented using a recursion operator. Secondly, we can re-use well-known results of the theory of automata (e.g., we used notions of trace equivalence and determinisation) and related tools.

Related work Automata-based models for specifying the local behaviour of distributed compontents are commonplace in the literature (see e.g., [13,21]). Less so is for the global specifications of choreographies: to the best of our knowledge, the conversation protocols (CP) in [26,27,9] (and references therein) are the only such model in the literature. The realisability of CP has been first studied in [27]; this is indeed the work closest to ours. Conversation protocols are non-deterministic Büchi automata whose labels resemble our interactions (barred the fact that, contrarily to our formalism, in [27] the sender and the receiver of each message are determined by its content). Our c-automata are basically finite-state automata where infinite words can be taken into account by looking at them as Büchi automata where all states are actually final. It is not immediate to provide a detailed comparison between conversation protocols and c-automata because their semantics and underlying communication models differ. As for the communication model, conversation protocols are realised in a subclass of CFSMs (cf. Section 5 of [27]), whereas we consider the unrestricted model of CFSMs, as well as a synchronous version of it. Concerning the semantics, Definition 4 (item 3(b)) of [27] restricts the runs to those where all messages in queues are eventually consumed, that is they require by definition a form of liveness. Instead, one of our goals is to identify conditions that guarantee relevant liveness properties. We prove them in Theorem 5.7, and in Prop. D.1 in the appendix we prove the exact property assumed in [27]. The realisability conditions of conversation protocols are *lossless join*, *synchronous compatibility*, and *autonomy*. Those conditions cannot easily be compared with well-formedness, due to the differences in the models and in the semantics. Furthermore, the style of the conditions is very different, and it also induces very different proof strategies in many cases. In particular,

– our well-sequencedness is checked on pairs of consecutive transitions and well-branchedness on pairs of coinitial paths;

– lossless join is a global property, that is a condition on the automaton consisting of the product of the languages of the local projections;
– synchronous compatibility is defined in terms of pairs of traces in the projection but verified with an algorithm that checks a global property of an automata costruction, and the same holds for autonomy.

Thus, while the conditions capture similar intuitions, a detailed comparison is very hard. When restricting to the common part of the two models, well-branchedness implies autonomy while the opposite does not hold. Indeed, by well-branchedness the selector is output-ready (according to the terminology in [27]), while any other participant either behaves uniformly in each branch (and is thus either input-ready or output-ready or termination-ready) or it is made aware of the choice by distinct inputs (that is it is input-ready). In all the cases autonomy is satisfied. In the other direction, a choice between traces $(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{l})(\mathsf{B}{\to}\mathsf{C}\colon \mathsf{n})(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{w})$ and $(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{r})(\mathsf{B}{\to}\mathsf{C}\colon \mathsf{n})(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{z})$ satisfies autonomy but not well-branchedness.

As for lossless join, we do not assume it. Actually, it is equivalent to one of our results, namely the correctness of the projection in the synchronous case (Thm. 4.14). Such a result is also used in the asynchronous case (Thm. 5.6), which is proved by reduction to the synchronous one via paired runs. We leave a detailed comparison of the two sets of constraints, in a common setting, for future work. Later works on CP (see, e.g., [9]) changed the approach and relied on model checking to show realisability instead of well-formedness conditions. Unfortunately, some of their main decidability results were flawed [24].

Conditions similar to well-branchedness and well-sequencedness do naturally arise in investigations related to choreoghraphies and their realisability. A unique sender driving a choice is a condition present in several multiparty session types formalisms ([28] and [20] to cite just a couple of them), global graphs formalisms [49], choreography languages in general (for instance see the notion of *dominant role* in [45]). Conditions related to item (3) of Def. 4.6 can also be found in multiparty session types formalisms [47] or in conversation protocols, as discussed above. Also, notions close to well-sequencedness turn out to arise quite naturally in "well-behaved" choreographies (see for instance the notion of *well-informedness* of [15] in the context of collaboration diagrams).

Similarly to what discussed in Remark 4.10, some approaches propose techniques to fix choreographies which are not well-behaved. This issue is considered in some multiparty session types [11,10], in algebraic and automata-based frameworks for choreographies [37,8] as well as in the choreographic middleware ChoreOS [3,4]. While they consider different conditions than ours, trying to adapt their approaches to our setting is an interesting item for future work.

As said, most approaches are not based on automata. For instance, [45,36,22] use algebraic operators to build larger choreographies from smaller ones, and give conditions on such operations ensuring that the resulting choreography is "well-behaved". This technique is not applicable in our case, since, like most works on automata, we do not consider an algebra to build automata.

While the main aim of c-automata is to provide a choreography model based on FSAs, we remark here that it is rather expressive and complements existing

models of choreographies or multiparty session types (MST). In particular, the expressive power of c-automata is not comparable with the one of the MST in [46], which subsumes most systems in the literature. More precisely, the c-automaton $\mathbb{C}$ref in § 1 cannot be syntactically written in [46] due to the two entangled loops. That example cannot be expressed in global graphs [49] either, again due to the intersecting loops. We note that the infinite unfolding of the c-automaton is regular and therefore it would fit in the session type system considered in [48]. However, this type system has not been conceived for choreographies (it is a binary session type system) and does not allow non-determinism.

On the other side, examples such as [46, Ex. 2, Fig. 4] cannot be written in our model (since we expect the same roles to occur in branches which are coinitial, branches inside loops require that all participants in a loop are notified when the loop ends). We conjecture that a refinement of well-branchedness is possible to address this limitation. Global graphs are another model of global specifications. Their advantage is that they feature parallel composition, which c-automata lack. We note however that one could use the classical product of automata on c-automata to model parallel composition in the case where the two branches have disjoint sets of participants (as typically assumed in MST with parallel composition). Mapping global graphs without parallel composition into c-automata is trivial. The same considerations apply to choreography languages where possible behaviours are defined by a suitable process algebra with parallel composition such as [36,14].

Future work One of the main motivations to develop a choreography model based on automata was to lift the compositional mechanism discovered in [6] on CFSMs to global specifications, in such a way that composition of global specifications preserves well-formedness. This is the problem we are currently addressing.

An interesting future development is also to adopt Büchi automata as c-automata. This extension is technically straightforward (just add accepting states to Def. 3.1 and define ω-languages accordingly), but it probably impacts greatly the underlying theory. An interesting yet not trivial effort is the identification of well-formedness conditions on this generalised class of c-automata that guarantee a precise correspondence with the ω-languages of the projections.

The interplay between FSAs and formal languages could lead to a theory of projection of choreographies based on languages instead of automata. For instance, one could try to characterise the languages accepted by well-formed c-automata, similarly to what done in [1,39,49]. In those approaches global specifications are rendered as partial orders and the distributed realisability is characterised in terms of closure properties of languages.

A final direction for future work concerns the implementation of tool support for the approach. We are currently working in this direction. A very preliminary and partial implementation by Simone Orlando and Ivan Lanese is available at `https://github.com/simoneorlando/Corinne`.

References

1. Rajeev Alur, Kousha Etessami, and Mihalis Yannakakis. Inference of Message Sequence Charts. *IEEE Trans. Software Eng.*, 29(7):623–633, 2003.
2. W. Ariola and C. Dunlop. Testing in the API Economy. Top 5 Myths. `https://api2cart.com/api-technology/api-testing-myths-infographic/`.
3. Marco Autili, Paola Inverardi, and Massimo Tivoli. Choreography realizability enforcement through the automatic synthesis of distributed coordination delegates. *Sci. Comput. Program.*, 160:3–29, 2018.
4. Marco Autili, Davide Di Ruscio, Amleto Di Salle, and Alexander Perucci. Choreosynt: enforcing choreography realizability in the future internet. In Shing-Chi Cheung, Alessandro Orso, and Margaret-Anne D. Storey, editors, *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, (FSE-22), Hong Kong, China, November 16 - 22, 2014*, pages 723–726. ACM, 2014.
5. Marco Autili, Amleto Di Salle, Francesco Gallo, Claudio Pompilio, and Massimo Tivoli. Chorevolution: Automating the realization of highly-collaborative distributed applications. In *COORDINATION*, pages 92–108, 2019.
6. Franco Barbanera, Ugo de'Liguoro, and Rolf Hennicker. Connecting open systems of communicating finite state machines. *J. Log. Algebraic Methods Program.*, 109, 2019.
7. Samik Basu and Tevfik Bultan. Automated choreography repair. In Perdita Stevens and Andrzej Wasowski, editors, *FASE*, volume 9633 of *Lecture Notes in Computer Science*, pages 13–30. Springer, 2016.
8. Samik Basu, Tevfik Bultan, and Meriem Ouederni. Deciding choreography realizability. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, pages 191–202, 2012.
9. Laura Bocchi, Julien Lange, and Emilio Tuosto. Amending contracts for choreographies. In *ICE*, volume 59 of *EPTCS*, pages 111–129, 2011.
10. Laura Bocchi, Julien Lange, and Emilio Tuosto. Three algorithms and a methodology for amending contracts for choreographies. *Sci. Ann. Comp. Sci.*, 22(1):61–104, 2012.
11. Jonas Bonér. *Reactive Microsystems - The Evolution Of Microservices At Scale*. O'Reilly, 2018.
12. Daniel Brand and Pitro Zafiropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, 1983.
13. Mario Bravetti and Gianluigi Zavattaro. Towards a unifying theory for choreography conformance and contract compliance. In Markus Lumpe and Wim Vanderperren, editors, *Software Composition, 6th International Symposium, SC 2007, Braga, Portugal, March 24-25, 2007, Revised Selected Papers*, volume 4829 of *Lecture Notes in Computer Science*, pages 34–50. Springer, 2007.
14. Tevfik Bultan and Xiang Fu. Specification of realizable service conversations using collaboration diagrams. *Service Oriented Computing and Applications*, 2(1):27–39, 2008.
15. Tevfik Bultan, Xiang Fu, Richard Hull, and Jianwen Su. Conversation specification: a new approach to design and analysis of e-service composition. In Gusztáv Hencsey, Bebo White, Yih-Farn Robin Chen, László Kovács, and Steve Lawrence, editors, *Proceedings of the Twelfth International World Wide Web Conference, WWW 2003, Budapest, Hungary, May 20-24, 2003*, pages 403–410. ACM, 2003.

16. Marco Carbone, Kohei Honda, and Nobuko Yoshida. Structured communication-centred programming for web services. In Rocco De Nicola, editor, *ESOP*, volume 4421 of *Lecture Notes in Computer Science*, pages 2–17. Springer, 2007.
17. Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In Roberto Giacobazzi and Radhia Cousot, editors, *POPL*, pages 263–274. ACM, 2013.
18. Gérard Cécé and Alain Finkel. Verification of programs with half-duplex communication. *I&C*, 202(2):166–190, 2005.
19. Mario Coppo, Mariangiola Dezani-Ciancaglini, Nobuko Yoshida, and Luca Padovani. Global progress for dynamically interleaved multiparty sessions. *Mathematical Structures in Computer Science*, 26(2):238–302, 2016.
20. Luca de Alfaro and Thomas A. Henzinger. Interface automata. In *Proceedings of the 8th European Software Engineering Conference held jointly with 9th ACM SIGSOFT International Symposium on Foundations of Software Engineering 2001, Vienna, Austria, September 10-14, 2001*, pages 109–120, 2001.
21. Pierre-Malo Deniélou and Nobuko Yoshida. Dynamic multirole session types. In Thomas Ball and Mooly Sagiv, editors, *POPL*, pages 435–446. ACM, 2011.
22. Pierre-Malo Deniélou and Nobuko Yoshida. Multiparty session types meet communicating automata. In Helmut Seidl, editor, *Programming Languages and Systems - 21st European Symposium on Programming, ESOP 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings*, volume 7211 of *Lecture Notes in Computer Science*, pages 194–213. Springer, 2012.
23. Alain Finkel and Étienne Lozes. Synchronizability of communicating finite state machines is not decidable. In *ICALP*, pages 122:1–122:14, 2017.
24. Adrian Francalanza, Claudio Antares Mezzina, and Emilio Tuosto. Reversible choreographies via monitoring in erlang. In Silvia Bonomi and Etienne Rivière, editors, *Distributed Applications and Interoperable Systems - 18th IFIP WG 6.1 International Conference, DAIS 2018, Held as Part of the 13th International Federated Conference on Distributed Computing Techniques, DisCoTec 2018, Madrid, Spain, June 18-21, 2018, Proceedings*, volume 10853 of *Lecture Notes in Computer Science*, pages 75–92. Springer, 2018.
25. Xiang Fu, Tevfik Bultan, and Jianwen Su. Conversation protocols: A formalism for specification and verification of reactive electronic services. In Oscar H. Ibarra and Zhe Dang, editors, *Implementation and Application of Automata, 8th International Conference, CIAA 2003, Santa Barbara, California, USA, July 16-18, 2003, Proceedings*, volume 2759 of *Lecture Notes in Computer Science*, pages 188–200. Springer, 2003.
26. Xiang Fu, Tevfik Bultan, and Jianwen Su. Conversation protocols: a formalism for specification and verification of reactive electronic services. *Theor. Comput. Sci.*, 328(1-2):19–37, 2004.
27. Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In George C. Necula and Philip Wadler, editors, *POPL*, pages 273–284. ACM Press, 2008.
28. Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. *Journal of the ACM*, 63(1):9:1–9:67, 2016. Extended version of a paper presented at POPL08.
29. Hans Hüttel, Ivan Lanese, Vasco T. Vasconcelos, Luís Caires, Marco Carbone, Pierre-Malo Deniélou, Dimitris Mostrous, Luca Padovani, António Ravara, Emilio Tuosto, Hugo Torres Vieira, and Gianluigi Zavattaro. Foundations of session types and behavioural contracts. *ACM Comput. Surv.*, 49(1):3:1–3:36, 2016.

30. Nickolas Kavantzas, David Burdett, Gregory Ritzinger, Tony Fletcher, Yves Lafon, and Charlton Barreto. Web services choreography description language version 1.0. Technical report, W3C, 2005. `http://www.w3.org/TR/ws-cdl-10/`.
31. Naoki Kobayashi. A Partially Deadlock-Free Typed Process Calculus. *ACM TOPLAS*, 20(2):436–482, 1998.
32. Naoki Kobayashi. Type-Based Information Flow Analysis for the Pi-Calculus. *Acta Informatica*, 42(4–5):291–347, 2005.
33. Dexter Kozen and Alexandra Silva. Practical Coinduction. *Mathematical Structures in Computer Science*, 27(7):1132–1152, 2017.
34. Ajay Krishna, Pascal Poizat, and Gwen Salaün. Checking business process evolution. *Sci. Comput. Program.*, 170:1–26, 2019.
35. Ivan Lanese, Claudio Guidi, Fabrizio Montesi, and Gianluigi Zavattaro. Bridging the gap between interaction- and process-oriented choreographies. In *Software Engineering and Formal Methods, SEFM 2008*, pages 323–332, 2008.
36. Ivan Lanese, Fabrizio Montesi, and Gianluigi Zavattaro. Amending choreographies. In António Ravara and Josep Silva, editors, *WWV*, volume 123 of *EPTCS*, pages 34–48, 2013.
37. Julien Lange, Emilio Tuosto, and Nobuko Yoshida. From communicating machines to graphical choreographies. In Sriram K. Rajamani and David Walker, editors, *POPL*, pages 221–232. ACM, 2015.
38. Markus Lohrey. Safe Realizability of High-Level Message Sequence Charts. In Luboš Brim, Mojmír Křetínský, Antonín Kučera, and Petr Jančar, editors, *CONCUR*, LNCS, pages 177–192. Springer, 2002.
39. Fabrizio Montesi. *Choreographic Programming*. PhD thesis, University of Copenhagen, 2013.
40. OMG. *Business Process Model and Notation (BPMN), Version 2.0*, January 2011. `https://www.omg.org/spec/BPMN`.
41. Luca Padovani, Vasco Thudichum Vasconcelos, and Hugo Torres Vieira. Typing liveness in multiparty communicating systems. In *COORDINATION*, volume 8459 of *LNCS*, pages 147–162. Springer, 2014.
42. Robert Paige and Robert Tarjan. Three Partition Refinement Algorithms. *SIAM Journal on Computing*, 16(6):973–989, December 1987.
43. Mila Dalla Preda, Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, and Jacopo Mauro. Dynamic choreographies: Theory and implementation. *Logical Methods in Computer Science*, 13(2), 2017.
44. Zongyan Qiu, Xiangpeng Zhao, Chao Cai, and Hongli Yang. Towards the theoretical foundation of choreography. In Carey L. Williamson, Mary Ellen Zurko, Peter F. Patel-Schneider, and Prashant J. Shenoy, editors, *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pages 973–982. ACM, 2007.
45. Alceste Scalas and Nobuko Yoshida. Less is more: multiparty session types revisited. *PACMPL*, 3(POPL):30:1–30:29, 2019.
46. Paula Severi and Mariangiola Dezani-Ciancaglini. Observational Equivalence for Multiparty Sessions. *Fundamenta Informaticae*, 170:267–305, 2019.
47. Paula Severi, Luca Padovani, Emilio Tuosto, and Mariangiola Dezani-Ciancaglini. On sessions and infinite data. *Logical Methods in Computer Science*, 13(2), 2017.
48. Emilio Tuosto and Roberto Guanciale. Semantics of global view of choreographies. *J. Log. Algebr. Meth. Program.*, 95:17–40, 2018.

A Proofs for Section 2 (Preliminaries)

Proposition 2.6. *Let* $\mathsf{S} = (M_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *and* $\mathsf{S}' = (M'_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *be two communicating systems. If* $\mathfrak{L}(M_{\mathsf{A}}) = \mathfrak{L}(M'_{\mathsf{A}})$ *for all* $\mathsf{A} \in \mathcal{P}$ *then* $\mathfrak{L}([\![\mathsf{S}]\!]^{\mathsf{s}}) = \mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{s}})$.

Proof. By contradiction, assume that there is $\sigma \in \mathfrak{L}([\![\mathsf{S}]\!]^{\mathsf{s}}) \setminus \mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{s}})$. Then there exists a finite prefix σ' and a suffix σ'' of σ such that

(a) $\sigma = \sigma' \cdot \sigma''$ and

(b) $\sigma' \in \mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{s}})$ and for any run $(\vec{q'}_{i-1} \xrightarrow{\lambda'_{i-1}} \vec{q'}_{i})_{1\leq i\leq n'}$ in S', if $\lambda'_0 \cdots \lambda'_h = \sigma'$ for some $0 \leq h < n'$ then no run in S' from $\vec{q'}_{h+1}$ accepts σ''.

Recall that it can be $n' = \infty$ and observe that $\sigma \neq \varepsilon$ (since both $\mathfrak{L}([\![\mathsf{S}]\!]^{\mathsf{s}})$ and $\mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{s}})$ accept ε) and $\sigma'' \neq \varepsilon$ (otherwise S' accepts $\sigma' \cdot \sigma'' = \sigma' = \sigma$).
We notice that the set

$$\{\, \mathit{length}(\delta) \mid \delta \textit{ is a prefix of } \sigma \textit{ satisfying conditions (a) and (b)} \,\}$$

is bounded, otherwise it would follow that $\sigma \in \mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{s}})$ (see Def. 2.2 and the observation following it). So, let us take as σ' the longest of the prefixes satisfying conditions (a) and (b) and let

$$(\vec{q}_{i-1} \xrightarrow{\lambda_{i-1}} \vec{q}_i)_{1\leq i\leq n} \tag{2}$$

be a run in S accepting σ.

We now study the behaviour of the CFSMs of S' after prefix σ'. For each participant A, consider the run $(\vec{q}_{i-1}(\mathsf{A}) \xrightarrow{\lambda_{i-1}\downarrow_{\mathsf{A}}} \vec{q}_i(\mathsf{A}))_{1\leq i\leq m}$ of M_{A} induced by the run in (4) obtained by removing transitions of the form $q \xrightarrow{\varepsilon} q$ from $(\vec{q}_{i-1}(\mathsf{A}) \xrightarrow{\lambda_{i-1}\downarrow_{\mathsf{A}}} \vec{q}_i(\mathsf{A}))_{1\leq i\leq n}$.

Assume $\sigma'' \in (\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}) \cdot \sigma'''$ for some σ''' (recall that $\sigma'' \neq \varepsilon$). Take any run

$$(\vec{q'}_{i-1} \xrightarrow{\lambda'_{i-1}} \vec{q'}_{i})_{1\leq i\leq n'} \tag{3}$$

in S' satisfying the conditions (a) and (b).

Observe that $(\lambda'_1 \cdots \lambda'_h)\downarrow_{\mathsf{A}} = \sigma'\downarrow_{\mathsf{A}} = (\lambda_1 \cdots \lambda_k)\downarrow_{\mathsf{A}}$ for some $0 \leq k < n$ and note that M_{A} can accept $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ from $\vec{q}_{k+1}(\mathsf{A})$. Using the same reasoning, M_{B} can accept $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ from a state reachable after performing $(\lambda'_1 \cdots \lambda'_h)\downarrow_{\mathsf{B}}$.

By condition (b) and the hypothesis on the length of σ', we have that $\vec{q'}_{h+1} \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}}\!\!\!\!\!\!/\;\;$. This implies that either M'_{A} cannot accept $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ from $\vec{q'}_{h+1}(\mathsf{A})$ or M'_{B} cannot accept $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ from $\vec{q'}_{h+1}(\mathsf{B})$. Then, contrary to the hypothesis, either M_{A} is not equivalent to M'_{A} or M_{B} is not equivalent to M'_{B}. □

The proof that the asynchronous semantics is invariant under language equivalence of CFSMs is very similar to the one of Prop. 2.8.

Proposition 2.8. *Let* $\mathsf{S} = (M_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *and* $\mathsf{S}' = (M'_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *be two communicating systems. If* $\mathfrak{L}(M_{\mathsf{A}}) = \mathfrak{L}(M'_{\mathsf{A}})$ *for all* $\mathsf{A} \in \mathcal{P}$ *then* $\mathfrak{L}([\![\mathsf{S}]\!]^{\mathsf{a}}) = \mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{a}})$.

Proof. By contradiction, assume that there is $\sigma \in \mathfrak{L}([\![\mathsf{S}]\!]^{\mathsf{a}}) \setminus \mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{a}})$. Then we can find a finite prefix σ' and a suffix σ'' of σ such that

(a) $\sigma = \sigma' \cdot \sigma''$ and

(b) for any run $(\vec{q'}_{i-1} \xrightarrow{\lambda'_{i-1}} \vec{q'}_i)_{1 \le i \le n'}$ in S', if $\lambda'_0 \cdots \lambda'_h = \sigma'$ for some $0 \le h < n'$ then no run in S' from $\vec{q'}_{h+1}$ accepts σ''.

Recall that it can be $n' = \infty$ and observe that $\sigma \neq \varepsilon$ (since both $\mathfrak{L}([\![\mathsf{S}]\!]^{\mathsf{a}})$ and $\mathfrak{L}([\![\mathsf{S}']\!]^{\mathsf{a}})$ accept ε) and $\sigma'' \neq \varepsilon$ (otherwise S' accepts $\sigma' \cdot \sigma'' = \sigma' = \sigma$).

Take as σ' the longest of the prefixes satisfying conditions (a) and (b) and let

$$(\vec{q}_{i-1} \xrightarrow{\lambda_{i-1}} \vec{q}_i)_{1 \le i \le n} \tag{4}$$

be a run in S accepting σ.

We now study the behaviour of the CFSMs of S' after prefix σ'. Let $\pi : \mathcal{L}_{\mathrm{act}} \times \mathcal{P} \to \mathcal{L}_{\mathrm{act}} \cup \{\varepsilon\}$ be the function defined as

$$\pi(\lambda, \mathsf{A}) = \begin{cases} \lambda, & \text{if } \lambda = \mathsf{A}\,\mathsf{B}!\mathsf{m} \text{ or } \lambda = \mathsf{B}\,\mathsf{A}?\mathsf{m} \\ \varepsilon, & \text{otherwise} \end{cases}$$

For each participant A, consider the run $(\vec{q}_{i-1}(\mathsf{A}) \xrightarrow{\pi(\lambda_{i-1},\mathsf{A})} \vec{q}_i(\mathsf{A}))_{1 \le i \le m}$ of M_{A} induced by the run in (4) obtained by removing transitions of the form $q \xrightarrow{\varepsilon} q$ from $(\vec{q}_{i-1}(\mathsf{A}) \xrightarrow{\pi(\lambda_{i-1},\mathsf{A})} \vec{q}_i(\mathsf{A}))_{1 \le i \le n}$.

We have that $\sigma'' \in (\lambda) \cdot \sigma'''$ for some σ''' and $\lambda \in \mathcal{L}_{\mathrm{act}}$ with subject say A (recall that $\sigma'' \neq \varepsilon$). Take any run

$$(\vec{q'}_{i-1} \xrightarrow{\lambda'_{i-1}} \vec{q'}_i)_{1 \le i \le n'} \tag{5}$$

in S' satisfying the conditions (a) and (b).

Observe that $(\lambda'_1 \cdots \lambda'_h)\!\downarrow_{\mathsf{A}} = \sigma'\!\downarrow_{\mathsf{A}} = (\lambda_1 \cdots \lambda_k)\!\downarrow_{\mathsf{A}}$ for some $0 \le k < n$ and note that M_{A} can accept λ from $\vec{q}_{k+1}(\mathsf{A})$. By condition (b) and the hypothesis on the length of σ', we have that $\vec{q'}_{h+1} \not\xrightarrow{\lambda}$. This implies that M'_{A} cannot accept λ from $\vec{q'}_{h+1}(\mathsf{A})$. Hence, contrary to the hypothesis, M_{A} is not equivalent to M'_{A}. □

B Proofs for Section 3 (Choreography Automata)

Proposition 3.5. *For all c-automata* $\mathbb{C}\mathrm{A}$ *and* $\mathsf{A} \in \mathcal{P}$, $\mathfrak{L}(\mathbb{C}\mathrm{A}\!\downarrow_{\mathsf{A}}) = \mathfrak{L}(\mathbb{C}\mathrm{A})\!\downarrow_{\mathsf{A}}$.

Proof. Let A_{A} be the intermediate automaton built in the first step in the projection of $\mathbb{C}\mathrm{A}$ on A. We have

$$\mathfrak{L}(\mathbb{C}\mathrm{A}\!\downarrow_{\mathsf{A}}) = \mathfrak{L}(A_{\mathsf{A}}) = \mathfrak{L}(\mathbb{C}\mathrm{A})\!\downarrow_{\mathsf{A}}$$

where the first equality follows since $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$ and A_{A} are trace equivalent by construction, and the second equality follows since A_{A} is obtained transition per transition by projecting the labels of $\mathbb{C}\mathrm{A}$ on A. □

Proposition 3.6. *If* $\mathbb{C}\mathrm{A}$ *and* $\mathbb{C}\mathrm{A}'$ *are trace-equivalent c-automata then* $\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{A}}}$ *and* $\mathbb{C}\mathrm{A}'{\downarrow_{\mathsf{A}}}$ *are isomorphic for each participant* $\mathsf{A} \in \mathcal{P}$.

Proof. Let us consider the intermediate automata A_{A} and A'_{A} in the projection of $\mathbb{C}\mathrm{A}$ and $\mathbb{C}\mathrm{A}'$, respectively. A_{A} and A'_{A} are trace equivalent, that is $\mathfrak{L}(A_{\mathsf{A}}) = \mathfrak{L}(A'_{\mathsf{A}})$, since by hypothesis $\mathfrak{L}(\mathbb{C}\mathrm{A}) = \mathfrak{L}(\mathbb{C}\mathrm{A}')$, $\mathfrak{L}(A_{\mathsf{A}}) = \mathfrak{L}(\mathbb{C}\mathrm{A}){\downarrow_{\mathsf{A}}}$ and $\mathfrak{L}(A'_{\mathsf{A}}) = \mathfrak{L}(\mathbb{C}\mathrm{A}'){\downarrow_{\mathsf{A}}}$ by Prop. 3.5 and since projection is a function. Given that A_{A} and A'_{A} are trace equivalent then their minimal deterministic automata with respect to trace equivalence are isomorphic. □

C Proofs for Section 4 (Well-formed Choreography Automata)

Remark C.1. It is worth noticing that, despite Thm. 2.6, the determinisation of intermediate automata is a necessary step in Def. 3.3 in order to obtain "well-behaved" projected systems. In fact, let us consider the following well-formed c-automaton.

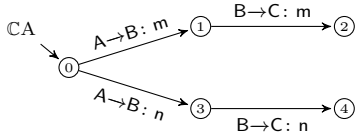

If we took out the determinisation step in Def. 3.3, we would obtain the following projected system.

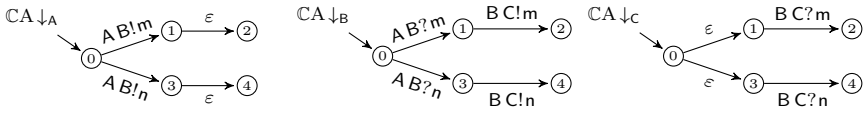

Now, the following one is a run of the above system.

$$(0,0,0) \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} (1,1,0) \xrightarrow{\varepsilon} (1,1,3)$$

Since in $\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{C}}}$ we have $3 \xrightarrow{\mathsf{B}\,\mathsf{C}?\mathsf{n}} 4$, but no run out of $(1,1,3)$ contains this transition as component transition, the system is not live. Besides, since we have also that $(1,1,3) \xrightarrow{\varepsilon} (2,1,3) \not\to$, configuration $(1,1,3)$ is a lock, and configuration $(2,1,3)$ is a deadlock.

Proposition 4.12. *Let* $\mathbb{C}\mathrm{A}$ *be a well-sequenced c-automaton. Then* $\mathfrak{L}(\mathbb{C}\mathrm{A})$ *is concurrency closed, i.e.* $\mathfrak{L}(\mathbb{C}\mathrm{A}) = \mathit{close}(\mathfrak{L}(\mathbb{C}\mathrm{A}))$.

Proof. Let us assume $\mathfrak{L}(\mathbb{C}\mathrm{A}) \neq \mathit{close}(\mathfrak{L}(\mathbb{C}\mathrm{A}))$. Clearly, $\mathfrak{L}(\mathbb{C}\mathrm{A}) \subseteq \mathit{close}(\mathfrak{L}(\mathbb{C}\mathrm{A}))$. Then there exists $w' \in \mathit{close}(\mathfrak{L}(\mathbb{C}\mathrm{A}))$ such that $w' \notin \mathfrak{L}(\mathbb{C}\mathrm{A})$. By definition of closure, there is $w \in \mathfrak{L}(\mathbb{C}\mathrm{A})$ such that $w \sim w'$. By definition of $\sim$, there exist two choreographic words w_1 and w_2 such that $w \sim w_1 \cdot (\mathsf{A}\to\mathsf{B}\colon \mathsf{b}) \cdot (\mathsf{C}\to\mathsf{D}\colon \mathsf{c}) \cdot w_2 \in \mathfrak{L}(\mathbb{C}\mathrm{A})$ and $w_1 \cdot (\mathsf{C}\to\mathsf{D}\colon \mathsf{c}) \cdot (\mathsf{A}\to\mathsf{B}\colon \mathsf{b}) \cdot w_2 \notin \mathfrak{L}(\mathbb{C}\mathrm{A})$, where $\{\,\mathsf{A},\mathsf{B}\,\} \cap \{\,\mathsf{C},\mathsf{D}\,\} = \emptyset$.

This implies that $\mathbb{C}$A is not well-sequenced. In fact, if that were not the case, let q and q' be the states of $\mathbb{C}$A from which the subwords $(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{b})\cdot(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{c})$ and w_2 are, respectively, generated. By well-sequencedness, there is also a path from q to q' generating $(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{c})\cdot(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{b})$, so implying that $w_1\cdot(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{c})\cdot(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{b})\cdot w_2 \in \mathfrak{L}(\mathbb{C}\mathrm{A})$, which is a contradiction. □

The proof of Thm. 4.14 and of the related lemmas requires to reason coinductively to account for possibly infinite runs. In particular, we adopt the coinduction style advocated in [34] which, without any loss of formal rigour, promotes readability and conciseness.

We begin with a technical lemma.

Lemma C.2. *Let* $\mathbb{C}$A *be a well-formed c-automaton and let* $q_1 \xrightarrow{\varepsilon} q_2 \xrightarrow{\varepsilon} \dots q_{n_{\mathsf{A}}-1} \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}} q_{n_{\mathsf{A}}}$ *in* A_{A}
and
$q_1 \xrightarrow{\varepsilon} q'_2 \xrightarrow{\varepsilon} \dots q'_{n_{\mathsf{B}}-1} \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}} q'_{n_{\mathsf{B}}}$ *in* A_{B}
where A_{A} *and* A_{B} *are the intermediate automata for* A *and* B *as defined in Def. 3.3. Then a transition with label* $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ *is enabled in the state* q_1 *of* $\mathbb{C}$A.

Proof. Consider the two runs above for A and B. We have two cases, depending on whether the two runs are projection of the same path in $\mathbb{C}$A or not. In the first case, the last transition is concurrent to all the previous ones. Indeed, the previous transitions are projected to ε both on A and on B, hence neither A nor B can occur in the label. Thus, thanks to well-sequencedness, a transition with the same label than the last one is enabled in the initial state as desired.

In the second case, let us assume that the two runs above are projections of different paths π_{A} and π_{B}. Note that π_{A} cannot contain occurrences of A and π_{B} cannot contain occurrences of B (but for the last transition). If one of the paths does not contain occurrences of neither A nor B the same reasoning above applies. Otherwise, π_{B} contains an occurrence of A. We can extend the two paths to obtain a span, hence thanks to well-branchedness the first different characters involving A should be different inputs. However, this is impossible since the first action involving A in π_{A} is the output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$. Also note that π_{B} cannot perform $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ since its transitions (but the last one, which is $\mathsf{A}\,\mathsf{B}?\mathsf{m}$) cannot involve B.

Now, a transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ is enabled in the state q_1 of $\mathbb{C}$A.

Before proving a further technical lemma, we observe that in case a transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ is enabled in the initial state of a c-automaton $\mathbb{C}$A, the projection of the c-automaton $\mathbb{C}\mathrm{A}'$ obtained by executing such a transition does not coincide with the system obtained by executing the synchronous transition in $\mathbb{C}\mathrm{A}{\downarrow}$, as shown by the following example.

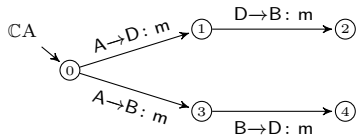

The projection of the c-automaton $\mathbb{C}\mathrm{A}'$ obtained by executing the transition labelled $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ is the system

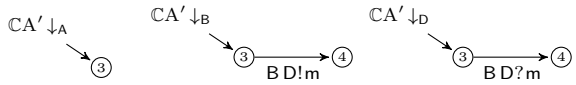

whereas the system obtained by executing the synchronous transition in $\mathbb{C}\mathrm{A}{\downarrow}$ is

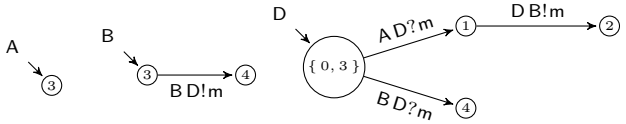

However, it is possible to show that the synchronous semantics of $\mathbb{C}\mathrm{A}'{\downarrow}$ is equivalent to the one $\mathbb{C}\mathrm{A}{\downarrow}$ after the transition $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$. In particular,

Lemma C.3. *Let* $\mathbb{C}\mathrm{A} = \langle Q, q_0, \mathcal{L}_{int}, \to\rangle$ *such that* $q_0 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q'$*, and let* $[\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}} = \langle \mathbb{S}, \vec{\mathbf{q_0}}, \mathcal{L}_{int}, \to_{[\![\,]\!]^{\mathsf{s}}}\rangle$ *such that* $\vec{\mathbf{q_0}} \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}}_{[\![\,]\!]^{\mathsf{s}}} \vec{\mathbf{q}}'$.
Then $\langle Q, q', \mathcal{L}_{int}, \to\rangle$ *and* $\langle \mathbb{S}, \vec{\mathbf{q}}', \mathcal{L}_{int}, \to_{[\![\,]\!]^{\mathsf{s}}}\rangle$ *are bisimilar.*

Proof. In order to avoid cumbersome notations, we often use the same symbol '$\to$' for the various transition relations of the automata we deal with, when the intended relation is clear from the context.
We first notice that, for each for $\mathsf{X} \in \mathcal{P} \setminus \{\mathsf{A}, \mathsf{B}\}$, we have $\vec{\mathbf{q}}'(\mathsf{X}) = \mathbf{q_{0_X}}$ where $\mathbf{q_{0_X}}$ is the initial state of $\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{X}}}$; moreover, $\vec{\mathbf{q_0}}(\mathsf{A}) \xrightarrow{\mathsf{A\,B!m}} \vec{\mathbf{q}}'(\mathsf{A})$ and $\vec{\mathbf{q_0}}(\mathsf{B}) \xrightarrow{\mathsf{A\,B?m}} \vec{\mathbf{q}}'(\mathsf{B})$. Now, if $\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{X}}} = \langle \mathbb{S}_{\mathsf{X}}, \mathbf{q_{0_X}}, \mathcal{L}_{\mathrm{act}}, \to^{\mathsf{X}}\rangle$ for $\mathsf{X} \in \mathcal{P}$, and $\mathbb{C}\mathrm{A}' = \langle Q, q', \mathcal{L}_{\mathrm{int}}, \to\rangle$, the thesis can be obtained as a consequence of the following statements:

i) $\mathbb{C}\mathrm{A}'{\downarrow_{\mathsf{X}}} = \langle \mathbb{S}_{\mathsf{X}}, \vec{\mathbf{q}}'(\mathsf{X}), \mathcal{L}_{\mathrm{act}}, \to_{\mathsf{X}}\rangle$ for $\mathsf{X} \in \{\mathsf{A}, \mathsf{B}\}$.
ii) $\mathfrak{L}(\mathbb{C}\mathrm{A}'{\downarrow_{\mathsf{X}}}) = \mathfrak{L}(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{X}}}) \setminus U$ for each $\mathsf{X} \in \mathcal{P} \setminus \{\mathsf{A}, \mathsf{B}\}$
where
$U = \bigcup_{\mathsf{n}\in\mathcal{M},\mathsf{Z}\in\mathcal{P}} \big(\{\, \mathsf{Z\,X?n} \cdot w \mid \mathsf{Z}{\to}\mathsf{X}\colon \mathsf{n} \notin \mathsf{hd}(\mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}})), w \in \mathcal{L}^*_{\mathrm{act}} \,\} \cup$
$\{\, \mathsf{X\,Z!n} \cdot w \mid \mathsf{X}{\to}\mathsf{Z}\colon \mathsf{n} \notin \mathsf{hd}(\mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}})), w \in \mathcal{L}^*_{\mathrm{act}} \,\}\big)$

Let us proceed with statement i).
Let us consider the case $\mathsf{X} = \mathsf{A}$ (the case $\mathsf{X} = \mathsf{B}$ being similar). We recall that, by definition of projection, we apply the minimisation algorithm on the intermediate automaton A_{A} as defined in Def. 3.3. In such an algorithm, an operation performing the transitive closure of the ε-transitions is used. We denote by ε-$\mathsf{clos}(\cdot)$ such an operation. Taking into account how the minimization algorithm works, and in particular the use of ε-$\mathsf{clos}(\cdot)$ in the definition of the transition function of the minimised automaton, it is enough to show that, for any $\hat{q} \in \varepsilon\text{-}\mathsf{clos}(q_0) = \vec{\mathbf{q_0}}(\mathsf{A})$ such that $\hat{q} \xrightarrow{\mathsf{A\,B!m}} q''$ (in A_{A}) we have that $q'' \in \varepsilon\text{-}\mathsf{clos}(q')$. So, let $\hat{q}$ be as above. By definition of ε-$\mathsf{clos}(\cdot)$, we have that there exists a run in $\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{A}}}$

$$\pi : q_0 \xrightarrow{\varepsilon}{}^{*} \hat{q} \xrightarrow{\mathsf{A\,B!m}} q''$$

We now show, by induction on the number n of the ε-transitions in π, that there exists also a number of ε-transitions from q' to q'', that is $q'' \in \varepsilon\text{-}\mathsf{clos}(q')$. If $n = 0$ there is nothing to prove. For $n > 0$, let the first transition in π be the projection of the transition $q_0 \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{n}} q_1$ of $\mathbb{C}\mathrm{A}$, where $\mathsf{A} \notin \{\mathsf{C}, \mathsf{D}\}$. By well-branchedness, we have also that $q' \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{n}} q_1'$ for a certain q_1', and $q_1 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q_1'$. Hence we have $q' \xrightarrow{\varepsilon} q_1'$ and $q_1 \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}} q_1'$ in A_{A}. The thesis then follows by the induction hypothesis.

Now, in order to prove the statement ii), let $\mathsf{X} = \mathsf{D} \in \mathcal{P} \setminus \{\mathsf{A}, \mathsf{B}\}$. We distinguish the following two possible cases:

$\mathsf{Z} = \mathsf{C} \notin \{\mathsf{A}, \mathsf{B}\}$.
Let us then take a string in $\mathfrak{L}(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{X}}})$ beginning with $\mathsf{D}\,\mathsf{C}!\mathsf{n}$ (the case of an input label can be treated similarly) and such that $\mathsf{D}\to\mathsf{C}\colon \mathsf{n} \in \mathsf{hd}(\mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}}))$. By , there exists a transition in $\mathbb{C}\mathrm{A}$ of the form $q_0 \xrightarrow{\mathsf{D}\to\mathsf{C}\colon \mathsf{n}} \tilde{q}$. Now, since with $\{\mathsf{C}, \mathsf{D}\} \cap \{\mathsf{A}, \mathsf{B}\} = \emptyset$ and since in $\mathbb{C}\mathrm{A}$ we have $q_0 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q'$, by well-sequencedness it follows that there exists q'' such that $q' \xrightarrow{\mathsf{D}\to\mathsf{C}\colon \mathsf{n}} q''$ and $\tilde{q} \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q''$. From $q' \xrightarrow{\mathsf{D}\to\mathsf{C}\colon \mathsf{n}} \tilde{q}''$ it follows that $\mathsf{D}\,\mathsf{C}!\mathsf{n}$ is also the label of a transition in $\mathbb{C}\mathrm{A}'{\downarrow_{\mathsf{D}}}$. Since $\tilde{q} \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q''$, the thesis follows then by coinduction.

$\mathsf{Z} \in \{\mathsf{A}, \mathsf{B}\}$. The cases when a string in $\mathfrak{L}(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{D}}})$ begins with $\mathsf{D}\,\mathsf{Z}!\mathsf{n}$ can be treated as in the previous item, using well-branchedness.
Let us then consider a string in $\mathfrak{L}(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{D}}})$ begins with $\mathsf{Z}\,\mathsf{D}?\mathsf{n}$ and such that $\mathsf{Z}\to\mathsf{D}\colon \mathsf{n} \in \mathsf{hd}(\mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}}))$.
We first observe that $q' \in \varepsilon\text{-}\mathsf{clos}(q_0)$ in A_{D}, so, by $\mathsf{Z}\to\mathsf{D}\colon \mathsf{n} \in \mathsf{hd}(\mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}}))$, using Lemma Appendix C, we have two consider two further subcases: either
1. $q_0 \xrightarrow{\mathsf{Z}\to\mathsf{D}\colon \mathsf{n}} \tilde{q}$ and $q' \xrightarrow{\mathsf{Z}\to\mathsf{D}\colon \mathsf{n}} q''$, or
2. $q' \xrightarrow{\mathsf{Z}\to\mathsf{D}\colon \mathsf{n}} q''$ and there exists no $\tilde{q}$ such that $q_0 \xrightarrow{\mathsf{Z}\to\mathsf{D}\colon \mathsf{n}} \tilde{q}$.
In the second case the thesis is immediate, in the second one it follows by the fact that $\mathfrak{L}(\mathbb{C}\mathrm{A}{\downarrow_{\mathsf{D}}})$ is deterministic. □

Theorem 4.14. $\mathfrak{L}(\mathbb{C}\mathrm{A}) = \mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}})$ *for any well-formed c-automaton* $\mathbb{C}\mathrm{A}$.

Proof. We first show that $\mathfrak{L}(\mathbb{C}\mathrm{A}) \subseteq \mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}})$.
Let us take a trace $\sigma \in \mathfrak{L}(\mathbb{C}\mathrm{A})$. We need to show that $\sigma \in \mathfrak{L}([\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}})$.

The proof is by coinduction. If σ is empty, there is nothing to prove.
Otherwise, let σ be $(\mathsf{A}\to\mathsf{B}\colon \mathsf{m})\cdot\sigma'$, where $\mathsf{A}\to\mathsf{B}\colon \mathsf{m}$ is accepted through a transition $q_0 \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} q'$ and where σ' is accepted by $\mathbb{C}\mathrm{A}$ starting from q'.
Consider now the intermediate automata. By construction, $q_0 \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}} q'$ in A_{A}, $q_0 \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}} q'$ in A_{B} and $q_0 \xrightarrow{\varepsilon} q'$ in A_{C} for each $\mathsf{C} \neq \mathsf{A}, \mathsf{B}$. This means that, by definition of synchronous semantics, $\mathsf{A}\to\mathsf{B}\colon \mathsf{m}$ is accepted by $[\![\mathbb{C}\mathrm{A}{\downarrow}]\!]^{\mathsf{s}}$ through a run of the form

$$\vec{q_0} \xrightarrow{\varepsilon} \vec{q_1} \xrightarrow{\varepsilon} \dots \vec{\tilde{q}} \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} \vec{q'}$$

where

1. the number of ε-transitions is equal to the number of participants in $\mathbb{C}$A minus two;
2. $\vec{\hat{q}}(\mathsf{C}) = q'$ for each $\mathsf{C} \notin \{\mathsf{A}, \mathsf{B}\}$ and $\vec{\hat{q}}(\mathsf{A}) = \vec{\hat{q}}(\mathsf{B}) = q_0$;
3. $\vec{q'}(\mathsf{C}) = q'$ for each $\mathsf{C} \in \mathcal{P}$.

By coinduction, σ' is accepted by $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}$ starting from the configuration $\vec{q'}$, and the thesis follows.

We now show that $\mathfrak{L}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}) \subseteq \mathfrak{L}(\mathbb{C}\mathrm{A})$. By Thm. 2.6 we can take into account as CMSMs the intermediate automata as defined in Def. 3.3.

Let σ be $(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}) \cdot \sigma'$, where $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ is accepted by a run

$$\vec{q}_0 \xrightarrow{\varepsilon} \vec{q}_1 \xrightarrow{\varepsilon} \dots \vec{q}_{n-1} \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} \vec{q}_n$$

Hence, by definition of synchronous semantics, there are runs
$\vec{q}_0(\mathsf{A}) \xrightarrow{\varepsilon} \vec{q}_1(\mathsf{A}) \xrightarrow{\varepsilon} \dots \vec{q}_{n_{\mathsf{A}}-1}(\mathsf{A}) \xrightarrow{\mathsf{A}\,\mathsf{B}!\mathsf{m}} \vec{q}_{n_{\mathsf{A}}}(\mathsf{A})$
$\vec{q}_0(\mathsf{B}) \xrightarrow{\varepsilon} \vec{q}_1(\mathsf{B}) \xrightarrow{\varepsilon} \dots \vec{q}_{n_{\mathsf{B}}-1}(\mathsf{B}) \xrightarrow{\mathsf{A}\,\mathsf{B}?\mathsf{m}} \vec{q}_{n_{\mathsf{B}}}(\mathsf{B})$
and
$\vec{q}_0(\mathsf{C}) \xrightarrow{\varepsilon} \vec{q}_1(\mathsf{C}) \xrightarrow{\varepsilon} \dots \vec{q}_{n_{\mathsf{C}}-1}(\mathsf{C}) \xrightarrow{\varepsilon} \vec{q}_{n_{\mathsf{C}}}(\mathsf{C})$ for each $\mathsf{C} \in \mathcal{P}$ such that $\mathsf{C} \notin \{\mathsf{A}, \mathsf{B}\}$.

For the sake of readability, let us define
$\hat{q}_{\mathsf{X}} = \vec{q}_{n_{\mathsf{X}}-1}(\mathsf{X})$, $q'_{\mathsf{X}} = \vec{q}_{n_{\mathsf{X}}}(\mathsf{X})$ for every $\mathsf{X} \in \mathcal{P}$.

We now show that a transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ is enabled in the initial state of $\mathbb{C}$A.

Consider the two runs above for A and B. Then, by Appendix C a transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ is enabled in the initial state of $\mathbb{C}$A. Consider the c-automaton obtained by executing such a transition

We can obtain the thesis by coinduction by showing that the projection of such an automaton coincides with the system obtained by executing the synchronous transition in $\mathbb{C}\mathrm{A}\downarrow$. This last statement can be obtained as a consequence of Lemma C.3. Notice that such a lemma exploits the fact that all components in the projection are deterministic. □

Theorem 4.15. *Given a well-formed c-automaton* $\mathbb{C}$A*, its projection* $(\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}})_{\mathsf{A}\in\mathcal{P}}$ *is live, lock-free, and deadlock-free with respect to the synchronous semantics.*

Thm. 4.15 is proved by Props. C.4, C.5, and C.6 below. We first show that the projection of a well-formed c-automaton is live namely, if there is a participant willing to take a transition, there is a computation executing such transition. (Recall that projected CFSMs are deterministic.)

Proposition C.4 Liveness. *Let* $\mathbb{C}$A *be a well-formed c-automaton. Then* $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}$ *is live.*

Proof. By Def. 2.9(i), we have to show that for each $\vec{q} \in \mathcal{R}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}})$, if there is $\mathsf{A} \in \mathcal{P}$ with an outgoing transition t from $\vec{q}(\mathsf{A})$ in $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$, then there exists a run of $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}$ including a transition which has t as component transition. Note that $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$ is deterministic hence the transition is not ε.

Since $\vec{q}$ is reachable, there is a run to $\vec{q}$ producing some word w. Consider the state q of $\mathbb{C}$A reached by accepting w (we can assume $\mathbb{C}$A deterministic). The c-automaton $\mathbb{C}\mathrm{A}_q$ obtained from $\mathbb{C}$A by setting q as initial state is well-formed. Now, assume there exists A such that $\vec{q}(\mathsf{A})$ has an outgoing transition t. Thanks to the definition of projection, there is a transition t' in $\mathbb{C}$A involving A and another participant, say B. Such a transition is on a run with word w' which has w as prefix. A run with word w' exists also in $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}$, and thanks to Thm. 4.14, it goes through $\vec{q}$, since the state reached is fully determined by the label (remember that CFSMs are deterministic), and it executes t as desired. This proves the thesis. □

Prop. C.5 below shows that the projection of a well-formed c-automaton is lock-free, namely that participants with enabled transitions can sooner or later fire one of them.

Proposition C.5 Lock freedom. *Let* $\mathbb{C}$A *be a well-formed c-automaton. Then* $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}$ *is lock-free.*

Proof. By Def. 2.9(ii) we have to show that for each $\vec{q} \in \mathcal{R}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}})$, $\vec{q}$ is not a lock. Let then $\vec{q} \in \mathcal{R}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}})$. Since $\vec{q}$ is reachable, there is a run to $\vec{q}$ producing some word w. Consider the unique state q of $\mathbb{C}$A obtained by accepting w (we can assume $\mathbb{C}$A deterministic). The c-automaton $\mathbb{C}\mathrm{A}_q$ obtained from $\mathbb{C}$A by setting q as initial state is well-formed. Now, assume there exists A such that $\vec{q}(\mathsf{A})$ has an outgoing transition t in $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$. Note that $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$ is deterministic hence the transition is not ε.

Thanks to the definition of projection, there is a transition t' in $\mathbb{C}$A involving A and another participant, say B. Transition t' is the last one of a run with word w' in $\mathbb{C}\mathrm{A}_q$ which has w as prefix. Thanks to Thm. 4.14 there is a transition t'' in $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}}$ with the same label, hence also involving A. The run leading to it it goes through $\vec{q}$, since the state reached is fully determined by the label (remember that CFSMs are deterministic), hence the transition is also a transition of $[\![\mathbb{C}\mathrm{A}_q\downarrow]\!]^{\mathsf{s}}$. We have to show that each run in $[\![\mathbb{C}\mathrm{A}_q\downarrow]\!]^{\mathsf{s}}$ has a transition involving A. This will imply the thesis.

Assume towards a contradiction that there is a run maximal w.r.t. prefix order not involving A. Such a run and the one ending with t'' are coinitial. They are projections of two paths in $\mathbb{C}$A. By removing possible common transitions we can produce two paths which are coinitial, say in q'. In case they have common nodes, we can produce two paths that are also cofinal, and removing possible duplicates of cycles we can hence obtain a q'-span. In case they do not have common nodes, we can obtain out of them a q'-span by possibly extending the path ending with t'' and extending or cutting the other one. Thanks to well-branchedness, since one of the two paths contains a transition with A, the other one also needs to include a transition with A, against the hypothesis. This proves the thesis. □

The next proposition shows that the communicating system obtained by projecting a well-formed c-automaton is deadlock-free.

Proposition C.6 Deadlock freedom. *If* $\mathbb{C}\mathrm{A}$ *is a well-formed c-automaton then* $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{s}}$ *is deadlock-free.*

Proof. By Def. 2.9(iii), we have to show that for any $\vec{q} \in \mathcal{R}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{s}})$, $\vec{q}$ is not a deadlock.
We prove the contrapositive. If participant $\mathsf{A} \in \mathcal{P}$ has an outgoing transition t from $\vec{q}(\mathsf{A})$ in $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$, then $\vec{q}$ has at least one outgoing transition. This follows from Prop. C.5. □

D Proofs for Section 5 (Asynchronous Communications)

Lemma 5.4. *Let* $\mathbb{C}\mathrm{A}$ *be a c-automaton. Then* $\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}) = \mathit{close}_{\mathrm{a}}(\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}))$.

Proof. Trivially, $\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}) \subseteq \mathit{close}_{\mathrm{a}}(\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}))$. Hence, we need to show the other inclusion. Any word in $\mathit{close}_{\mathrm{a}}(\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}))$ is obtained from some word w of $\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}})$ by a finite number of swaps as in Def. 5.3. Hence, if we show that $\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}})$ is closed under single swaps then the thesis will follow by induction.

Let $w\,(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m})\,(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{n})\,w' \in \mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}})$ with $\mathsf{C} \notin \{\mathsf{A}, \mathsf{B}\}$ and let $\pi = s \xrightarrow{\mathsf{A}\to\mathsf{B}\colon \mathsf{m}} s_1 \xrightarrow{\varepsilon} \ldots \xrightarrow{\varepsilon} s_n \xrightarrow{\mathsf{C}\to\mathsf{D}\colon \mathsf{n}} s'$ be the portion of the accepting sender trace of $\mathbb{C}\mathrm{A}$ corresponding to the labels $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ and $\mathsf{C}{\to}\mathsf{D}\colon \mathsf{n}$. The thesis follows by exhibiting an accepting trace for $w\,(\mathsf{C}{\to}\mathsf{D}\colon \mathsf{n})\,(\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m})\,w'$.

Since $\mathsf{C}{\to}\mathsf{D}\colon \mathsf{n}$ is enabled in $s_n = \langle \vec{q_n} \;;\; \vec{b_n} \rangle$ then $\mathsf{C}\,\mathsf{D}!\mathsf{n}$ is enabled in $\vec{q_n}(\mathsf{C})$. Note that each ε-transition on π is an input since CFSMs are deterministic and thanks to the definition of sender traces. If none of the ε-transitions on π corresponds to a receive for C then the transition with label $\mathsf{C}\,\mathsf{D}!\mathsf{n}$ is enabled since the beginning, hence one can build a new trace by anticipating its execution. Otherwise one can anticipate both all the inputs for C and the transition $\mathsf{C}\,\mathsf{D}!\mathsf{n}$ preserving their relative order. Indeed, inputs can be anticipated since none of the transitions before them produces the message that they consume (since all these transitions are inputs). In either cases we have the thesis. □

Proposition D.1. *Let* $\mathbb{C}\mathrm{A}$ *be an asynchronously well-formed c-automaton. Each run* σ *in* $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}$ *is pairable. Furthermore, the run* σ' *into which it is paired is obtained from* σ *by anticipating input transitions.*

Proof. By coinduction. Let σ be a run in $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathrm{a}}$. Since in the initial configuration s_0 the queues are empty the first transition in σ is some output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$.

We now show that after $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ we can immediately execute the corresponding input $\mathsf{A}\,\mathsf{B}?\mathsf{m}$.

Since $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ is enabled there is a path in $\mathbb{C}\mathrm{A}$ composed by transitions where A does not occur followed by a transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$. By asynchronous well-sequencedness there are commuting diamonds with sides $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ and all previous transitions, hence there is a transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ from the initial state q_0.

By definition of projection there is a transition $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ enabled in the initial state of B, and from the definition of asynchronous semantics we know that the

message for B is available in the queue. Hence, the desired input transition with label A B?m can be taken.

We now have to show that there is a run σ' starting with A B!m followed by A B?m producing the same sender trace as σ.

Consider the first transition for B in σ. There are three cases:

B takes no transition in σ: note that after executing A B!m as discussed above a transition with label A B?m is enabled, and it will remain enabled unless B executes a transition. Hence, this case is disallowed since it would violate fairness.

the first transition taken by B in σ has label A B?m: since B is deterministic, this leads B in the same state as in σ', hence σ' can be obtained from σ by anticipating the transition with label A B?m. Note that this can be done since previous transitions do not involve B. Also, since we are just moving an input the sender trace is unchanged. This concludes the proof for this case.

the first transition taken by B in σ does not have label A B?m: the transition with label A B?m and the other one are generated from two different paths in $\mathbb{C}$A. The two paths are coinitial, and by dropping redundant parts we can transform them into a span. Participant A is a sender in q_0, hence it is also a sender in all the transitions in q_0 involving A, and such transitions are different. If the first transition t in σ involves A then the first different characters involving B should be different inputs, but since the first message in the queue is from A B!m then A B?m should trigger first, against the hypothesis. Hence this case can never happen. Thus, the first transition in σ can not involve A. By well-branchedness they need to be concurrent, hence there is a commuting diamond between the two transitions. Also, the sender cannot be B, otherwise B will not be the sender in all transitions involving it, violating well-branchedness. The same reasoning repeats for all transitions till A→B: m is finally reached. Hence, a transition with label A→B: m is enabled in the initial state of $\mathbb{C}$A.

Consider now the run σ. Since the input from A is enabled, thanks to fairness it should trigger in a finite number of steps. We show that we can anticipate it, thus building a suitable σ'. If the transition before it does not involve B at all, then the two transitions are independent and we can anticipate the input. If the transition sends a message to B, again they can be swapped, since they consider different queues. If the transition is an input from B then such an input belongs to a transition concurrent to A→B: m, hence by asynchronous well-sequencedness B has also a path taking the inputs in the opposite order, and we can then anticipate the second input as well. Thus σ' can be built from σ just by anticipating the desired input. Note that in all the cases the generated sender trace is unchanged, hence the thesis follows.

□

Theorem 5.6. *Let* $\mathbb{C}$A *be an asynchronously well-formed c-automaton.*

$$\mathfrak{L}^{!}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}}) = \mathfrak{L}(\mathbb{C}\mathrm{A})$$

Proof. Thanks to Prop. D.1 $\mathfrak{L}^!([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}})$ coincides with the language generated by pairable runs. Note that a pairable run corresponds to a synchronous run (indeed executing an output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ followed by the corresponding input $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ corresponds to execute the synchronous transition with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$), and the sender trace of the former coincides with the trace of the latter, hence

$$\mathfrak{L}^!([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}}) = \mathfrak{L}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{s}})$$

Note that $\mathbb{C}\mathrm{A}$ is also synchronously well-formed (the notion of well-branchedness is the same while the notion of asynchronous well-sequencedness implies the one of synchronous well-sequencedness).

Then the thesis follows thanks to Thm. 4.14. □

Thm. 5.7 is proved by Props. D.2, D.3, and D.4 below.

Proposition D.2 Lock freedom. *Let $\mathbb{C}\mathrm{A}$ be an asynchronously well-formed c-automaton. Then $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}}$ is lock-free.*

Proof. By Def. 2.9(ii) we have to show that for each configuration $s \in \mathcal{R}([\![\mathsf{S}]\!]^{\mathsf{a}})$, if there is $\mathsf{A} \in \mathcal{P}$ with an outgoing transition t from $s(\mathsf{A})$ in M_{A}, then each run of $[\![\mathsf{S}]\!]^{\mathsf{a}}$ from s maximal w.r.t. prefix order has a transition from A. Note that $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$ is deterministic hence the transition is not ε.

If t is an output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$, there is trivially a run executing it, since outputs can always trigger. Also, the output remains enabled unless a transition from A triggers. If such a transition is executed the thesis follows, otherwise by fairness $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ is executed and the thesis follows as well.

Let us now consider the case where t is an input $\mathsf{A}\,\mathsf{B}?\mathsf{m}$. It is obtained as projection of some transition $\hat{t}$ with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ in the original c-automaton. The proof is by induction on the distance between the initial state of the c-automaton and $\hat{t}$.

In the base case the distance is 0. Thus, the input $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ is enabled in the initial state of the CFSM for B. Also, the corresponding output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ is enabled in the CFSM for A. Hence, by definition of asynchronous semantics there exists a run in which the output is executed, followed by the input. Let us consider any other run. The transition with label $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ is enabled, and the only way to disable it is by executing another transiton from A. If such a transition is executed the thesis follows, otherwise by fairness $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ is executed and the thesis follows as well.

Let us now move to the inductive case, where the distance between the initial state of the c-automaton and $\hat{t}$ is greater than 0. Consider any transition $\hat{t}'$ whose distance is 0. Thanks to fairness, the corresponding output should be in every run, and since there are no dependences blocking it, it can be anticipated so to become the first transition in the run. The corresponding input can be moved just afterwards thanks to Prop. D.1. Let σ be the resulting run. If the original maximal run has a state where an input from some participant C is enabled and no transition from C is executed in the run, then the same issue occurs in σ (since the only difference w.r.t. the original run is that input and output of a different

transition have been anticipated). However, σ without the two first transitions is also a run of the c-automaton obtained by moving the initial state one step forward. In such a case the distance between the $\hat{t}$ and the initial state is shorter, hence the issue cannot occur by inductive hypothesis. □

Proposition D.3 Liveness. *Let* $\mathbb{C}\mathrm{A}$ *be an asynchronously well-formed c-automaton. Then* $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}}$ *is live.*

Proof. By Def. 2.9(i) we have to show that for each configuration $s \in \mathcal{R}([\![\mathsf{S}]\!]^{\mathsf{a}})$, if there is $\mathsf{A} \in \mathcal{P}$ with an outgoing transition t from $s(\mathsf{A})$ in M_{A}, then there is a run of $[\![\mathsf{S}]\!]^{\mathsf{a}}$ from s including a transition which has t as component transition. Note that $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$ is deterministic hence the transition is not ε.

If t is an output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$, there is trivially a run executing it, since outputs can always trigger.

Let us now consider the case where t is an input $\mathsf{A}\,\mathsf{B}?\mathsf{m}$. It is obtained as projection of some transition $\hat{t}$ with label $\mathsf{A}{\to}\mathsf{B}\colon \mathsf{m}$ in the original c-automaton. The proof is by induction on the distance between the initial state of the c-automaton and $\hat{t}$.

In the base case the distance is 0. Thus, the input $\mathsf{A}\,\mathsf{B}?\mathsf{m}$ is enabled in the initial state of the CFSM for B. Also, the corresponding output $\mathsf{A}\,\mathsf{B}!\mathsf{m}$ is enabled in the CFSM for A. Hence, by definition of asynchronous semantics there exists a run in which the output is executed, followed by the input.

Let us now move to the inductive case, where the distance between the initial state of the c-automaton and $\hat{t}$ is greater than 0. Consider any transition $\hat{t}'$ whose distance is 0. Thanks to fairness, the corresponding output should be in every run, and since there are no dependences it can be anticipated so to become the first transition in the run. The corresponding input can be moved just afterwards thanks to Prop. D.1. Let σ be the resulting run. If the original run has a configuration s where an input from some participant C is enabled in $\vec{s}(C)$ and there is no continuation in which it is executed, then the same issue occurs in σ. However, σ without the two first transitions with all possible continuations are also runs of the c-automaton obtained by moving the initial state one step forward. In such a case the distance between $\hat{t}$ and the initial state is shorter, hence the issue cannot occur by inductive hypothesis. □

Proposition D.4 Deadlock freedom. *Let* $\mathbb{C}\mathrm{A}$ *be an asynchronously well-formed c-automaton. Then* $[\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}}$ *is deadlock-free.*

Proof. By Def. 2.9(iii), we have to show that for any $\vec{q} \in \mathcal{R}([\![\mathbb{C}\mathrm{A}\downarrow]\!]^{\mathsf{a}})$, $\vec{q}$ is not a deadlock.

We prove the contrapositive. If participant $\mathsf{A} \in \mathcal{P}$ has an outgoing transition t from $\vec{q}(\mathsf{A})$ in $\mathbb{C}\mathrm{A}\downarrow_{\mathsf{A}}$, then $\vec{q}$ has at least one outgoing transition. This follows from Prop. D.2. □